arXiv:1208.4863v2 [math.CO] 2 Oct 2012

Eigenvalues and Quasirandom Hypergraphs

John Lenz
University of Illinois at Chicago
lenz@math.uic.edu

Dhruv Mubayi *
University of Illinois at Chicago
mubayi@math.uic.edu

October 4, 2012

Abstract

Let $p(k)$ denote the partition function of k. For each $k \geq 2$, we describe a list of $p(k) - 1$ quasirandom properties that a k-uniform hypergraph can have. Our work connects previous notions on hypergraph quasirandomness, beginning with the early work of Chung and Graham and Frankl-Rödl related to strong hypergraph regularity, the spectral approach of Friedman-Wigderson, and more recent results of Kohayakawa-Rödl-Skokan and Conlon-Hàn-Person-Schacht on weak hypergraph regularity and its relation to counting linear hypergraphs.

For each of the quasirandom properties that are described, we define a hypergraph eigenvalue analogous to the graph case and a hypergraph extension of a graph cycle of even length whose count determines if a hypergraph satisfies the property. Our work can therefore be viewed as an extension to hypergraphs of the seminal results of Chung-Graham-Wilson for graphs.

Our results yield the following applications. First, motivated by Sidorenko's Conjecture on the minimum homomorphism density of bipartite graphs in arbitrary graphs, we show that an analog of the conjecture for hypergraphs holds for a variety of hypergraph cycles. These are the first infinite families of hypergraphs with minimum degree two where this has been verified. Second, we give an efficient certification algorithm for hypergraph quasirandomnes which leads to an efficient strong refutation algorithm for random k-SAT. For n-vertex, k-uniform hypergraphs with $k \geq 4$ and at least $n^{k/2+\sqrt{k}}$ edges, we provide an algorithm with running time $O(n^{k\omega} \operatorname{polylog} n)$ that certifies quasirandomness for almost all hypergraphs. This improves the previous best running time for such certification due to Coja-Oghlan-Cooper-Frieze and Hán-Person-Schacht, in addition to also certifying a stronger quasirandom property than these previous results.

*Research supported in part by NSF Grant 0969092.

2010 Mathematics Subject Classification. Primary 05C80; Secondary 05C50, 05C65.

Key words and phrases. quasirandom, hypergraph, eigenvalue.

1 Introduction

The study of quasirandom or pseudorandom graphs was initiated by Thomason [74, 75] where he studied deterministically constructed graph sequences $\mathcal{G} = \{G_n\}_{n\to\infty}$ that behaved similar to the Erdős-Rényi random graph $G(n,p)$. Thomason introduced the notion of jumbled graphs, which are graphs where the density of edges in every large subset of vertices does not deviate very far from the expected number of edges that the random graph would have. Thomason supplied several examples of explicitly constructed jumbled graphs, found a polynomial time checkable property for jumbled graphs, and investigated many of the properties of these graphs. In a subsequent paper, Chung, Graham, and Wilson [18], building on this and other work, provided a list of (deterministic) properties of graph sequences which were inspired by $G(n,p)$ and proved that these properties are equivalent. Beginning with these foundational papers on the subject [18, 74, 75], the last two decades have seen an explosive growth in the study of quasirandom structures in mathematics and computer science. Quasirandom structures lie at the heart of recent proofs of Szemerédi's Theorem [71] using the Strong Hypergraph Regularity Lemma [34, 56, 58, 59, 73]. They also have many applications in computer science including the algorithmic version of the Regularity Lemma [3], expander graphs [38], and recent work relating expander graphs to the PCP Theorem and hardness of approximation (see [26]). For details on quasirandomness, we refer the reader to a survey of Krivelevich and Sudakov [45] for graphs and recent papers of Gowers [33, 34, 35] for general quasirandom structures including hypergraphs.

1.1 Previous Results

The core of what Chung, Graham, and Wilson [18] proved is the following remarkable theorem. We write $o(n^2)$ for a function $f : \mathbb{N} \to \mathbb{R}$ such that $f(n)n^{-2} \to 0$ as $n \to \infty$. The precise way to interpret the $o(\cdot)$ terms in the following theorems is explained in the discussion after Theorem 4. As usual, for a graph or hypergraph G, we write $V(G)$ for its vertex set and $E(G)$ for its edge set. Unless specifically mentioned, all graphs and hypergraphs in this paper have no loops or multiple edges. For $U \subseteq V(G)$, the induced subgraph on U, denoted $G[U]$, is the graph with vertex set U and edge set $\{e \in E(G) : e \subseteq U\}$. If F and G are graphs or hypergraphs, a *labeled copy of F in G* is an edge-preserving injection $V(F) \to V(G)$, i.e. an injection $\alpha : V(F) \to V(G)$ such that if E is an edge of F, then $\{\alpha(x) : x \in E\}$ is an edge of G.

Theorem 1. ***(Chung-Graham-Wilson [18])*** *Let $0 < p < 1$ be a fixed constant and let $\mathcal{G} = \{G_n\}_{n\to\infty}$ be a sequence of graphs such that $|V(G_n)| = n$ and $|E(G_n)| \geq p\binom{n}{2} + o(n^2)$. The following properties are equivalent:*

- *`Disc`: (short for discrepency) For every $U \subseteq V(G_n)$, $|E(G_n[U])| = p\binom{|U|}{2} + o(n^2)$.*
- *`Expand`: For every $S, T \subseteq V(G_n)$, $e(S,T) = p|S||T| + o(n^2)$, where $e(S,T)$ is the number of edges with one endpoint in S and one endpoint in T, with edges inside $S \cap T$ counted twice.*

- *`Count[All]`: For every fixed graph F with f vertices and e edges, the number of labeled copies of F in G_n is $p^e n^f + o(n^f)$.*
- *`Cycle`$_4$: The number of labeled copies of C_4 in G_n is at most $p^4 n^4 + o(n^4)$.*
- *`Cycle`$_{2\ell}$: For $\ell \geq 2$ a fixed integer, the number of labeled copies of $C_{2\ell}$ in G_n is at most $p^{2\ell} n^{2\ell} + o(n^{2\ell})$.*
- *`Eig`: If λ_1 and λ_2 are the first and second largest eigenvalues in absolute value of the adjacency matrix of G_n respectively, then $|\lambda_1| = pn + o(n)$ and $|\lambda_2| = o(n)$.*

Chung, Graham, and Wilson [18] proved several more properties, but the five properties stated in Theorem 1 are usually viewed as central to the field. In particular, `Disc` and `Expand` are basic properties of random graphs. Note that Chung, Graham, and Wilson [18] put the condition "$|E(G_n)| \geq p\binom{n}{2} + o(n^2)$" into the statement of the properties, but since it appears in every property the condition can be moved to a requirement on the graph sequence as in Theorem 1.

Since Chung, Graham, and Wilson's [18] work, many more properties have been added into the equivalence in Theorem 1, see [39, 49, 54, 55, 57, 60, 61, 65, 66, 67, 68, 76]. In addition, several researchers investigated the sparse case where p is no longer a constant but $p = o(1)$, see [2, 11, 12, 42, 43]. In this paper, we will be concentrating only on the dense case when p is a fixed constant.

Almost immediately after proving Theorem 1, Chung and Graham [10, 13, 15, 16, 17] investigated generalizing the theorem to k-uniform hypergraphs. A *k-uniform hypergraph* is a pair of finite sets $(V(H), E(H))$ where $E(H)$ is a collection of k-subsets of $V(H)$. One initial difficulty in generalizing Theorem 1 to $k > 2$ is a construction of Erdős and Hajnal [27].

Construction. (Erdős and Hajnal [27]) Consider the three-uniform hypergraph formed as follows (actually, a sequence of probability spaces of n vertex hypergraphs, one for each integer n). Let T_n be a random (graph) tournament on n vertices, by which we mean an orientation of the edges of the complete graph K_n where each edge is oriented in one of two directions with equal probability and the orientations are done independently for distinct edges. Form an n-vertex, three-uniform hypergraph H_n by letting three vertices u, v, w form an edge of H_n if $u \to v \to w \to u$.

Rödl observed that this construction shows that the hypergraph generalizations of `Disc` and `Count[All]` are not equivalent. The sequence H_n is easily seen to satisfy `Disc`, since in any $U \subseteq V(H_n)$ the probability that three vertices from U form an edge of H_n is $p = \frac{1}{4}$, the same as the expected density of H_n. On the other hand, it is easily observed that H_n does not contain any copy of $K_4^{(3)}$, the complete three-uniform hypergraph on four vertices. Since the number of $K_4^{(3)}$ in H_n is zero, the graph sequence H_n obviously fails `Count[All]`. Motivated by this, Chung and Graham [13, 15, 16, 17] investigated how to strengthen the properties `Disc`/`Expand` to make them equivalent to `Count[All]`. They found several properties equivalent to `Count[All]`; the main property they use is related to the number of even/odd subgraphs of a given hypergraph which they called `Deviation`. Building on a

suggestion of Frankl and Rödl, they also showed that a strengthened version of `Disc` which we call `CliqueDisc` is equivalent to `Count[All]`. For a $(k-1)$-uniform hypergraph G, define the *cliques of order k in G*, denoted $\mathcal{K}_k(G)$, to be the collection of k-sets $T \subseteq V(G)$ such that all $(k-1)$-subsets of T are edges of G. If $\{H_n\}_{n\to\infty}$ is a hypergraph sequence, the property `CliqueDisc` states that for every $(k-1)$-uniform hypergraph G, the edges of H_n intersect a p-proportion of the cliques of order k in G. Note that for $k=2$ this is the same as `Disc`. The papers [13, 15, 16, 17] only focused on the case $p=\frac{1}{2}$. Kohayakawa, Rödl, and Skokan [44] extended this to the case of fixed $0<p<1$ using several new techniques. The following theorem is a combination of the main results of this series of papers.

Theorem 2. ***(Chung and Graham [13, 15, 16, 17] and Kohayakawa-Rödl-Skokan [44])*** *Let $0<p<1$ be a fixed constant and let $\mathcal{H}=\{H_n\}_{n\to\infty}$ be a sequence of k-uniform hypergraphs such that $|V(H_n)|=n$ and $|E(H_n)|=p\binom{n}{k}+o(n^k)$. The following properties are equivalent:*

- *`CliqueDisc`$[k-1]$: For every $(k-1)$-uniform hypergraph G with $V(G)=V(H_n)$, $|H_n\cap\mathcal{K}_k(G)|=p|\mathcal{K}_k(G)|+o(n^k)$.*
- *`Count[All]`: For every fixed k-uniform hypergraph F with e edges and f vertices, the number of labeled copies of F in H_n is $p^e n^f+o(n^f)$.*
- *`Count[Octahedron]`: The number of labeled copies of the k-octahedron O_k in H_n is $p^{2^k}n^{2k}+o(n^{2k})$, where the k-octahedron is the hypergraph with $2k$ vertices $x_{i,0},x_{i,1}$ for $1\le i\le k$ and edge set $\{\{x_{1,\epsilon_1},\dots,x_{k,\epsilon_k}\}:\epsilon_i\in\{0,1\}\}$.*

Theorem 2 only states some of the properties shown equivalent by Chung and Graham (see also [10]) and Kohawakaya-Rödl-Skokan [44]. They investigated several other properties including induced subhypergraph counting, common neighborhood sizes, and a property central to many of the proofs called `Deviation`. One cannot hope to divide a general hypergraph up into quasirandom pieces in the sense of Theorem 2. Indeed, this is shown by the construction of Rödl mentioned earlier. With an eye towards a proof of Szemerédi's Theorem [71], Gowers [33] investigated in more detail this obstacle and extended upon the quasirandom properties in Theorem 2. Eventually, the correct framework for partitioning a general hypergraph was stated in the Strong Hypergraph Regularity Lemma [34, 56, 58, 59, 73]. Keevash [40] extended on Gowers [33] work by proving an embedding type lemma for quasirandom hypergraphs where his notion of quasirandom was governed by `CliqueDisc` and applied it to some Turán type problems. Austin and Tau [8] build on and apply these quasirandom properties to investigate hypergraph property testing.

Theorem 2 and its generalizations provide several properties, chiefly `CliqueDisc`, that are equivalent to `Count[All]` in hypergraphs. However, it remained open whether the simpler property `Disc` for k-uniform hypergraphs is equivalent to counting some substructure like a generalization of C_4. Recently, Kohayakawa-Nagle-Rödl-Schacht [41] answered this affirmatively by showing that `Disc` is equivalent to counting linear hypergraphs, where a hypergraph H is *linear* if for every edge E in H, every other edge $E'\neq E$ intersects E in

at most one vertex. Building on this, Conlon-Hàn-Person-Schacht [22] showed that Disc is equivalent to counting a type of linear four cycle. These two results can be combined into the following theorem.

Theorem 3. ***(Kohayakawa-Nagle-Rödl-Schacht [41] and Conlon-Hàn-Person-Schacht [22])*** *Let $0 < p < 1$ be a fixed constant and let $\mathcal{H} = \{H_n\}_{n\to\infty}$ be a sequence of k-uniform hypergraphs such that $|V(H_n)| = n$ and $|E(H_n)| \geq p\binom{n}{k} + o(n^k)$. The following properties are equivalent:*

- *Disc: For every $U \subseteq V(H_n)$, $|E(H_n[U])| = p\binom{|U|}{k} + o(n^k)$.*
- *Count[linear]: For every fixed linear k-uniform hypergraph F with e edges and f vertices, the number of labeled copies of F in H_n is $p^e n^f + o(n^f)$.*
- *$\mathtt{Cycle}_4$: The number of labeled copies of C_4 in H_n is at most $p^{|E(C_4)|}n^{|V(C_4)|} + o(n^{|V(C_4)|})$, where C_4 is a linear hypergraph defined precisely in Section 2.*

Note that Conlon et al. [22] put the condition "$|E(H_n)| \geq p\binom{n}{k} + o(n^k)$" into the statement of the properties that don't trivially imply it like Disc and this is equivalent to the way we have stated Theorem 3. Conlon et al. [22] have several more properties including induced subgraph counts, and common neighborhood sizes, but we consider the properties stated in Theorem 3 as the core properties. Theorem 3 equates Disc, an exponential time checkable property, with $\mathtt{Cycle}_4$, a property which can be checked in polynomial time. Dellamonica and Rödl [25] add another property to Theorem 3 involving the count of linear hypergraphs in $H_n[X]$ for every vertex set X settling a conjecture of Conlon et al. [22] and Shapira and Yuster [62] add another property into Theorem 3 involving weighted hypergraph cuts.

1.2 Our Results

Chung and Graham and Kohayakawa-Rödl-Skokan did not provide a generalization of Eig into Theorem 2 and in addition Conlon et al. [22] did not provide a generalization of Eig to add to the equivalences of Theorem 3. For graphs, Eig helps enormously to simplify the proof that $\mathtt{Count}[C_4] \Rightarrow \mathtt{Expand}$. In addition, a spectral characterization of quasirandomness has proved to be a very useful result in many situations, for example in the proofs that certain explicitly constructed graphs are quasirandom (see [5, 6, 46, 70]).

This leads to our first main contribution. We define a generalization of Eig to k-uniform hypergraphs and add it into the equivalences stated in Theorem 3. This facilitates the proof of the implication $\mathtt{Cycle}_4 \Rightarrow \mathtt{Disc}$ and also aids the development of spectral methods of quasirandom hypergraphs, a direction that has proved very fruitful for graphs.

Our second main contribution is to generalize Theorem 3 substantially. Let $k \geq 2$ be an integer and let π be a proper partition of k, by which we mean that π is an unordered list of at least two positive integers whose sum is k. For the partition π of k given by $k = k_1 + \cdots + k_t$, we will abuse notation by saying that $\pi = k_1 + \cdots + k_t$. For every proper partition π, we define properties $\mathtt{Expand}[\pi]$, $\mathtt{Eig}[\pi]$, and $\mathtt{Cycle}_4[\pi]$ and show that they are equivalent.

Definition. *Let $k \geq 2$ and let $\pi = k_1 + \cdots + k_t$ be a proper partition of k. A k-uniform hypergraph F is* π-linear *if there exists an ordering $E_1, \ldots, E_m$ of the edges of F such that for every i, there exists a partition of the vertices of E_i into $A_{i,1}, \ldots, A_{i,t}$ such that for $1 \leq s \leq t$, $|A_{i,s}| = k_s$ and for every $j < i$, there exists an s such that $E_j \cap E_i \subseteq A_{i,s}$.*

The following is our main theorem.

Theorem 4. ***(Main Result)*** *Let $0 < p < 1$ be a fixed constant and let $\mathcal{H} = \{H_n\}_{n\to\infty}$ be a sequence of k-uniform hypergraphs such that $|V(H_n)| = n$ and $|E(H_n)| \geq p\binom{n}{k} + o(n^k)$. Let $\pi = k_1 + \cdots + k_t$ be a proper partition of k. The following properties are equivalent:*

- ***Eig**[π]: $\lambda_{1,\pi}(H_n) = pn^{k/2} + o(n^{k/2})$ and $\lambda_{2,\pi}(H_n) = o(n^{k/2})$, where $\lambda_{1,\pi}(H_n)$ and $\lambda_{2,\pi}(H_n)$ are the first and second largest eigenvalues of H_n with respect to π, which are defined in Section 3.*
- ***Expand**[π]: For all $S_i \subseteq \binom{V(H_n)}{k_i}$ where $1 \leq i \leq t$,*

$$e(S_1, \ldots, S_t) = p\prod_{i=1}^{t} |S_i| + o\left(n^k\right)$$

 where $e(S_1, \ldots, S_t)$ is the number of tuples $(s_1, \ldots, s_t)$ such that $s_1 \cup \cdots \cup s_t$ is a hyperedge and $s_i \in S_i$.

- ***Count**[π-linear]: If F is an f-vertex, m-edge, k-uniform, π-linear hypergraph, then the number of labeled copies of F in H_n is $p^m n^f + o(n^f)$.*
- ***Cycle**$_4$[π]: The number of labeled copies of $C_{\pi,4}$ in H_n is at most $p^{|E(C_{\pi,4})|}n^{|V(C_{\pi,4})|} + o(n^{|V(C_{\pi,4})|})$, where $C_{\pi,4}$ is the hypergraph four cycle of type π which is defined in Section 2.*
- ***Cycle**$_{4\ell}$[π]: the number of labeled copies of $C_{\pi,4\ell}$ in H_n is at most $p^{|E(C_{\pi,4\ell})|}n^{|V(C_{\pi,4\ell})|} + o(n^{|V(C_{\pi,4\ell})|})$, where $C_{\pi,4\ell}$ is the hypergraph cycle of type π and length 4ℓ defined in Section 2.*

Remarks.

- Theorem 4 is stated for hypergraph sequences $\mathcal{H} = \{H_n\}_{n\to\infty}$ where $|V(H_n)| = n$ for every n, but this requirement can be weakened. All the above theorems work perfectly well when considering sequences of hypergraphs whose number of vertices goes to infinity but perhaps might not be defined for every n. One way of enhancing the statements of these theorems to this situation is to transition away from the little-o notation and use an ϵ-δ type statement. Instead, following Chung, Graham, and Wilson [18], we maintain the language of sequences of hypergraphs as follows. Let $\mathcal{H} = \{H_{n_q}\}_{q\to\infty}$ be a sequence of hypergraphs such that $|V(H_{n_q})| = n_q$, $n_q < n_{q+1}$,

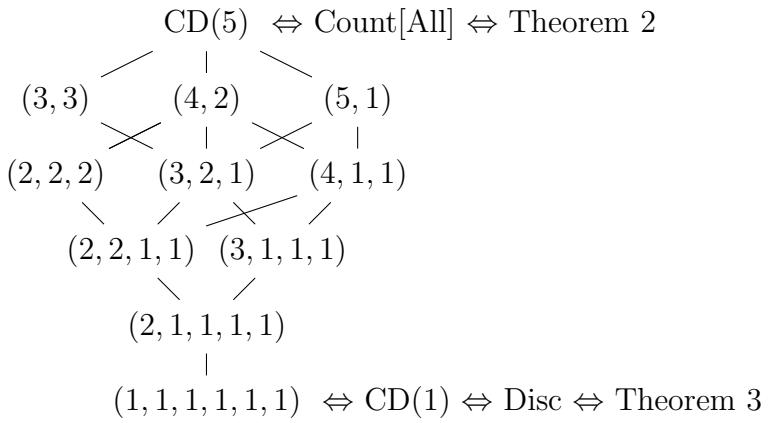


Figure 1: The Hasse diagram of quasirandom properties for $k = 6$.

and $|E(H_{n_q})| \geq p\binom{n_q}{k} + o(n_q^k)$, where now the little-$o$ expression means there exists a function $f(q)$ such that $|E(H_{n_q})| \geq p\binom{n_q}{k} + f(q)$ with $\lim_{q\to\infty} f(q)n_q^{-k} = 0$. Similarly, when we say that property P (which might include a little-o expression) implies a property P', what we mean is that there exist functions $f(q)$ and $f'(q)$ such that $P(f(q))$ implies $P'(f'(q))$, where the notation $P(f(q))$ stands for the property P with the little-o replaced by the function $f(q)$. In the rest of this paper, we ignore these complications and work directly with hypergraph sequences defined for every n but all the proofs extend into this world of hypergraph subsequences.

- If $\pi = 1 + \cdots + 1$, the partition of k into k ones, then the equivalences $\texttt{Expand}[\pi]$ $\Leftrightarrow$ $\texttt{Count}[\pi\text{-linear}]$ $\Leftrightarrow$ $\texttt{Cycle}_4[\pi]$ of Theorem 4 constitute Theorem 3. In this case, $\texttt{Expand}[1, \ldots, 1]$ is equivalent to $\texttt{Disc}$, π-linear is the same as a linear hypergraph, and the cycle $C_{(1,\ldots,1),4}$ we define in Section 2 is the same as the cycle defined by Conlon et al. [22].

- Similar to the way Chung, Graham, and Wilson [18] state Theorem 1, we could move the condition "$|E(H_n)| \geq p\binom{n}{k} + o(n^k)$" into the statement of the properties $\texttt{Count}[\pi$-linear$]$, $\texttt{Cycle}_{4\ell}[\pi]$, and $\texttt{Eig}[\pi]$, but since $\texttt{Expand}[\pi]$ trivially implies $|E(H_n)| \geq p\binom{n}{k} + o(n^k)$, moving the condition into the statement of the properties is equivalent to the way we have stated Theorem 4.

- With a little more work, we can extend Theorem 4 to include $\texttt{Cycle}_{2\ell}[\pi]$ in the case where $\pi = k_1 + k_2$ is a partition into two parts. That is, if $\pi = k_1 + k_2$, then having the correct count of any even length cycle implies all the conditions in Theorem 4. For partitions into more than two parts, we are unsure if counting cycles of lengths congruent to 2 modulo 4 can be added into Theorem 4. The simplest open case is counting six cycles with $\pi = 1 + 1 + 1$.

- If π' is a refinement of π, then clearly $\texttt{Count}[\pi\text{-linear}] \Rightarrow \texttt{Count}[\pi'\text{-linear}]$ and so if

$\{H_n\}_{n\to\infty}$ is a sequence satisfying the properties in Theorem 4 for π, it satisfies the properties for π'. In a companion paper [47], we show the converse: if π' is not a refinement of π then $\texttt{Expand}[\pi] \not\Rightarrow \texttt{Expand}[\pi']$ so the property $\texttt{Expand}[\pi]$ is distinct for distinct π and arranged in a poset via partition refinement. This is shown in Figure 1. In addition, [47] also compares $\texttt{Expand}[\pi]$ to essentially all previous hypergraph quasirandom properties and completely determines the poset of implications between these quasirandom properties. This provides the first step in a project recently suggested by Chung [10].

- Our proof of Theorem 4 works with ordered proper partitions $\vec{\pi} = (k_1, \ldots, k_t)$ because it is easier to work with the hypergraph properties in Theorem 4 when there is a fixed labeling of the parts of the partition π. It is easy to see that $\texttt{Expand}[\pi]$ and $\texttt{Count}[\pi$-linear] are the same property no matter which ordering is chosen for π. When defining $C_{\pi,2\ell}$ and $\lambda_{1,\pi}$ and $\lambda_{2,\pi}$, we will first define them for an ordering $\vec{\pi}$ of π and then observe that the definitions are independent of the ordering.

- For graphs, a family $\mathcal{F}$ of graphs is called *p-forcing* if $\forall F \in \mathcal{F}$, the number of labeled copies of F in G is $(1+o(1))p^{|E(F)|}n^{|V(F)|}$ implies that G is quasirandom in the sense of Theorem 1. Theorem 1 proves that $\{K_2, C_{2t}\}$ is p-forcing and Skokan and Thoma [68] proved that $\{K_2, K_{s,t}\}$ is p-forcing for $s, t \geq 2$. Skokan and Thoma [68] also posed the forcing conjecture: for every bipartite F which is not a forest, $\{K_2, F\}$ is p-forcing. This conjecture is still open. Dellamonica, Haxell, Łuczak, Mubayi, Nagle, Person, Rödl, and Schacht [24] investigated what properties of G are forced by $\{K_2, T\}$ when T is a tree, and Conlon, Fox, and Sudakov [21] proved the forcing conjecture holds for a class of bipartite graphs. Say that a family $\mathcal{F}$ of k-uniform hypergraphs is (π, p)-forcing if $\forall F \in \mathcal{F}$, the number of labeled copies of F in H is $(1+o(1))p^{|E(F)|}n^{|V(F)|}$ implies that H is π-quasirandom in the sense of Theorem 4. Theorem 4 proves that $\{K_k^{(k)}, C_{\pi,4\ell}\}$ is (π, p)-forcing. It would be interesting to determine which families are (π, p)-forcing.

1.3 Applications

Property $\texttt{Eig}[\pi]$, which we define and include in Theorem 4, can be used in a variety of settings. To show $\texttt{Eig}[\pi]$'s versatility, we give two applications of the hypergraph eigenvalue techniques which we develop in this paper.

- A beautiful conjecture of Sidorenko [63] asks about the minimum number of copies of a bipartite graph H which must be contained in another graph G. There is a simple upper bound on this minimum as follows. Let $G(n,p)$ be the random graph with edge density p; it is easy to see that w.h.p. if p is not too small then $G(n,p)$ has at most $p^{|E(H)|}n^{|V(H)|}$ labeled copies of H. Sidorenko's Conjecture [63] is that this bound is sharp for bipartite graphs. The formulation of this conjecture is usually in terms of *homomorphisms* or *edge-preserving maps*. Let $h_H(G)$ denote the number of labeled copies of H in G, i.e. the number of edge-preserving maps from $V(H)$ to $V(G)$. Let

$t_H(G) = h_H(G)/|V(G)|^{|V(H)|}$ be the normalization of $h_H(G)$, which is the fraction of mappings $f : V(H) \to V(G)$ which are edge preserving. Sidorenko's Conjecture is the following.

Conjecture 5. ***(Sidorenko [63])*** *For every bipartite graph H with m edges and every graph G,*

$$t_H(G) \geq t_{K_2}(G)^m.$$

Sidorenko's Conjecture is known to be true in a small number of cases, including complete bipartite graphs, trees, even cycles (see [63]), cubes [37], and a larger class including bipartite graphs with a vertex complete to the other part [21]. Lovász [48] proved a local version of the conjecture. For k-uniform hypergraphs with $k \geq 3$, Sidorenko [63] gave a construction (the loose triangle) for which the generalization of Conjecture 5 to k-uniform hypergraphs fails. Nevertheless, Sidorenko [64] proved that the generalization of Conjecture 5 holds for so-called "hypergraph trees": hypergraphs where the edges can be ordered so that there is a degree one vertex in the union of the first i edges in this ordering for every i. This leads to the following problem.

Problem 6. *For which k-partite, k-uniform hypergraphs H with m edges is it true that every k-uniform hypergraph G has $t_H(G) \geq t_{K_k}(G)^m$?*

An easy consequence of hypergraph eigenvalue techniques used to prove Theorem 4 is that this is true for all cycles of type π and length 4ℓ, for any π and any ℓ.

Theorem 7. *For every $k \geq 2$, every proper partition π of k, every $\ell \geq 1$, and every k-uniform hypergraph G,*

$$t_{C_{\pi,4\ell}}(G) \geq t_{K_k}(G)^{|E(C_{\pi,4\ell})|}.$$

The forcing conjecture mentioned previously is a stability result for Sidorenko's Conjecture, since the forcing conjecture asserts that if G is a graph which matches the bound in Conjecture 5, then G is structurally similar to $G(n,p)$, i.e. quasirandom.

- The k-SAT problem is one of the most studied NP-complete problems, and the hypergraph eigenvalue techniques developed for Theorem 4 give a simple polynomial-time algorithm to detect strong hypergraph quasirandomness, which leads to a polynomial-time strong refutation algorithm for random k-SAT. These strong refutation algorithms have been studied by several researchers [19, 20, 28, 31, 36]. It is easy to see that refuting random k-SAT can be reduced to verifying if a random k-uniform hypergraph is quasirandom as follows. Let ϕ be a random k-SAT formula with m clauses and n variables. Form a k-uniform hypergraph H on $2n$ vertices, one vertex for each variable and its negation. Each clause of ϕ produces a hyperedge of H, so that satisfying assignments of ϕ correspond to the complements of independent sets of size n in H. If

we can certify that H is quasirandom in the sense of `Expand`, H will have no large independent sets implying that ϕ has no satisfying assignment. The focus is determining the smallest m for which there exist such an algorithm.

Currently, the best result for $k \geq 4$ is that there exists a polynomial-time algorithm which refutes almost all k-SAT formulas if the number of clauses is at least $Cn^{\lceil k/2 \rceil}$ where n is the number of variables and C is a large constant. There are two algorithms that achieve this bound, one by Coja-Oghlan, Cooper, and Frieze [19] and one by Hán, Person, and Schacht [36]. The algorithm of Hán, Person, and Schacht [36] is based on Theorem 3 and [22], and works by counting $C_{1+\cdots+1,4}$.

As a consequence of the ideas developed for Theorem 4, we also provide an algorithm which counts cycles which certifies hypergraph quasirandomness in polynomial time. Compared to the algorithms of Coja-Oghlan, Cooper, and Frieze [19] and Hán, Person, and Schacht [36], our algorithm is significantly faster and certifies a stronger notion of quasirandomness at the expense of slightly weaker bounds on the number of clauses/number of hyperedges. Our algorithm runs in time $O(n^{k\omega} \operatorname{polylog} n)$ where ω is the exponent for matrix multiplication compared to the $O(n^{k2^{k-1}})$ running time of [36]. In addition, our algorithm certifies a version of $\texttt{Expand}[\pi]$ for $\pi = \lceil k/2 \rceil + \lfloor k/2 \rfloor$, while [36] certifies $\texttt{Expand}[1 + \cdots + 1]$ and [19] works with k-SAT formulas instead of hypergraphs. Our algorithm only works when the number of clauses is at least $n^{k/2+\epsilon}$ for even k and $n^{k/2+\sqrt{k}}$ for odd k, compared to the $Cn^{\lceil k/2 \rceil}$ of [19, 36].

Theorem 8. *Let $k \geq 4$ and $\epsilon, \delta > 0$ be given. There exists an algorithm which runs in time $O(n^{k\omega} \operatorname{polylog} n)$ where ω the exponent for matrix multiplication such that on input of an n-vertex, k-uniform hypergraph the algorithm outputs either `Quasirandom` or `Unknown` such that the following holds.*

- *If the algorithm outputs `Quasirandom` for an n-vertex, m-edge, k-uniform hypergraph H, then the following property holds: for all $S_1 \subseteq \binom{V(H)}{\lceil k/2 \rceil}$ and all $S_2 \subseteq \binom{V(H)}{\lfloor k/2 \rfloor}$,*

$$\left| e(S_1, S_2) - \frac{k!m}{n^k} |S_1||S_2| \right| \leq \frac{\delta m}{n^{k/2}} \sqrt{|S_1||S_2|}.$$

- *With high probability, the algorithm outputs `Quasirandom` on a hypergraph drawn from the distribution $G^{(k)}(n,p)$ when*

$$p \geq \begin{cases} n^{-k/2+\epsilon} & \text{if } k \text{ is even,} \\ n^{-k/2+\sqrt{k}} & \text{otherwise.} \end{cases}$$

The remainder of this paper is organized as follows. In Section 2, we define the hypergraph cycles $C_{\pi,4}$. Section 3 gives the formal definition of eigenvalues with respect to π. Theorem 4 is proved by showing a chain of implications in the order stated in the theorem; Section 4

proves $\texttt{Eig}[\pi] \Rightarrow \texttt{Expand}[\pi]$, Section 5 proves $\texttt{Expand}[\pi] \Rightarrow \texttt{Count}[\pi\text{-linear}]$, and Section 6 completes the proof of Theorem 4 by showing that $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$. Section 7 contains the proofs of Theorems 7 and 8. Lastly, Appendix A contains some linear algebra facts required for the proofs. Throughout this paper, we use the notation $[n] = \{1, \ldots, n\}$.

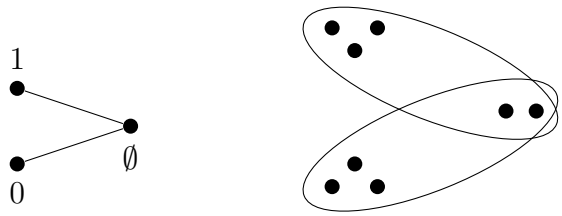
2 Hypergraph Cycles

In this section, we define the hypergraph cycles $C_{\pi,4}$ used in Theorem 4. The hypergraph cycles $C_{\pi,2\ell}$ are defined by first defining steps, then defining a path as a combination of steps, and finally defining the cycle as a path with its endpoints identified.

Definition. *Let $\vec{\pi} = (1, \ldots, 1)$ be the ordered partition of t into t parts. Define the* step of type $\vec{\pi}$*, denoted $S_{\vec{\pi}}$, as follows. Let A be a vertex set of size 2^{t-1} where elements are labeled by binary strings of length $t-1$ and let $B_2, \ldots, B_t$ be disjoint sets of size 2^{t-2} where elements are labeled by binary strings of length $t-2$. The vertex set of $S_{\vec{\pi}}$ is the disjoint union $A \dot\cup B_2 \dot\cup \cdots \dot\cup B_t$. Make $\{a, b_2, \ldots, b_t\}$ a hyperedge of $S_{\vec{\pi}}$ if $a \in A$, $b_j \in B_j$, and the code for b_{j+1} is equal to the code formed by removing the jth bit of the code for a.*

For a general $\vec{\pi} = (k_1, \ldots, k_t)$, start with $S_{(1,\ldots,1)}$ and enlarge each vertex into the appropriate size; that is, a vertex in A is expanded into k_1 vertices and each vertex in B_j is expanded into k_j vertices. More precisely, the vertex set of $S_{\vec{\pi}}$ is $(A \times [k_1]) \dot\cup (B_2 \times [k_2]) \dot\cup \cdots \dot\cup (B_t \times [k_t])$, and if $\{a, b_2, \ldots, b_t\}$ is an edge of $S_{(1,\ldots,1)}$, then $\{(a, j_1), (b_2, j_2), \ldots, (b_t, j_t)\}$ is a hyperedge of $S_{\vec{\pi}}$ for all $(j_1, \ldots, j_t) \in [k_1] \times \cdots \times [k_t]$.

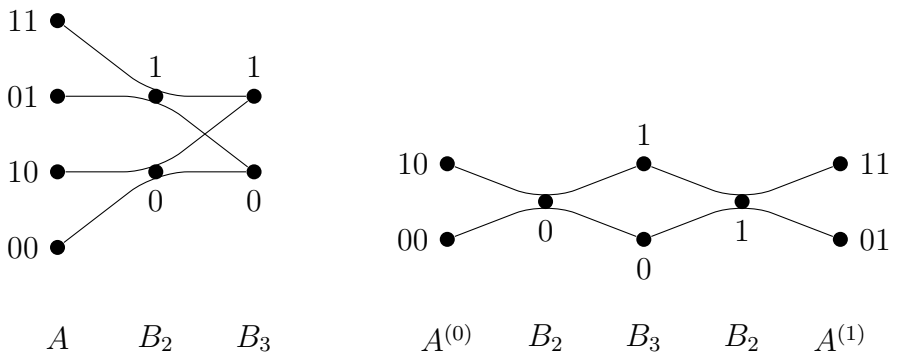
This defines the step of type $\vec{\pi}$*, denoted $S_{\vec{\pi}}$. Let $A^{(0)}$ be the ordered tuple of vertices of A in $S_{\vec{\pi}}$ whose binary code ends with zero and $A^{(1)}$ the ordered tuple of vertices of A whose binary code ends with one, where vertices are listed in lexicographic order within each $A^{(i)}$. These tuples $A^{(0)}$ and $A^{(1)}$ are the two* attach tuples of $S_{\vec{\pi}}$



(a) $S_{(1,1)}$ (b) $S_{(3,2)}$

Figure 2: Steps with $t = 2$

Figure 2 shows the steps of type $(1, 1)$ and type $(3, 2)$. Notice that each step has "length" two if we consider the attach tuples as the "ends" of a path.



(a) $S_{(1,1,1)}$ (b) $S_{(1,1,1)}$

Figure 3: Steps of type $\pi = (1, 1, 1)$

Figure 3 shows two different drawings of the step of type $\vec{\pi} = (1, 1, 1)$. Notice that the attach tuples are easily visible in Figure 3 *(b)*, since the two attach tuples are the codes in A ending with a zero and a one. The step of type $\vec{\pi} = (k_1, k_2, k_3)$ is an enlarged version of Figure 3 similar to Figure 2 *(b)*.

In general for arbitrary $\vec{\pi}$, the step $S_{\vec{\pi}}$ can be drawn in two ways similar to Figure 3. First from the definition, a step is a k-partite hypergraph with parts $A, B_2, \ldots, B_t$ so it can be drawn similar to Figure 3 *(a)*. But the step can also be drawn with the two attach tuples on separate ends of the picture like Figure 3 *(b)*. Let M_0 be the set of edges incident to vertices in the attach tuple $A^{(0)}$ and M_1 the set of edges incident to vertices in $A^{(1)}$. Edges from M_0 and M_1 intersect only in vertices in B_t because if $a_0 \in A^{(0)}$ and $a_1 \in A^{(1)}$ then the code for a_0 ends in a zero and the code for a_1 ends in a one, so only when deleting the last bit will the codes possibly be the same. Therefore, the step $S_{\vec{\pi}}$ can be viewed as a type of length two path in a hypergraph formed from a collection of k-partite edges M_0 between $A^{(0)}$ and B_t and another collection of k-partite edges M_1 between B_t and $A^{(1)}$. It is easy to see that all vertices in $\cup B_i$ have degree two and that the vertices in the attach tuples have degree one so are ready to be connected into longer paths.

Definition. *The* path of type $\vec{\pi}$ of length 2ℓ*, denoted $P_{\vec{\pi},2\ell}$, is the hypergraph formed from ℓ copies of $S_{\vec{\pi}}$ with successive attach tuples identified. That is, let $T_1, \ldots, T_\ell$ be copies of $S_{\vec{\pi}}$ and let $A_i^{(0)}$ and $A_i^{(1)}$ be the attach tuples of T_i. The hypergraph $P_{\vec{\pi},2\ell}$ is the hypergraph consisting of $T_1, \ldots, T_\ell$ where the vertices of $A_i^{(1)}$ are identified with $A_{i+1}^{(0)}$ for every $1 \leq i \leq \ell - 1$. (Recall that by definition, $A_i^{(1)}$ and $A_{i+1}^{(0)}$ are tuples (i.e. ordered lists) of vertices, so the identification of $A_i^{(1)}$ and $A_{i+1}^{(0)}$ identifies the corresponding vertices in these tuples.) The* attach tuples of $P_{\vec{\pi},2\ell}$ *are the tuples $A_1^{(0)}$ and $A_\ell^{(1)}$.*

In Figure 4, the path $P_{(1,1,1),4}$ is drawn as two copies of $S_{(1,1,1)}$ with attach tuples identified. The diamond, circle, and square vertices keep track of the parts A, B_2, B_3. For a general $P_{(k_1,k_2,k_3),4}$, each diamond vertex is enlarged into k_1 vertices, each circle vertex is enlarged

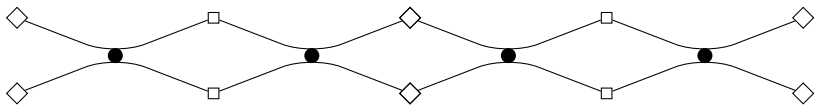


Figure 4: $P_{(1,1,1),4}$

into k_2 vertices, and each square vertex is enlarged into k_3 vertices. For a general $\vec{\pi}$, every step can be visualised as in Figure 3 *(b)* as two collections of k-partite edges M_0 and M_1 between A and B_t, so all paths $P_{\vec{\pi},2\ell}$ can be visualised as in Figure 4 as a concatenation of steps. Vertices in A have degree one and all other vertices have degree two, so the path has internal degree two vertices with degree one attach tuples.

Definition. *The* cycle of type $\vec{\pi}$ and length 2ℓ*, denoted* $C_{\vec{\pi},2\ell}$*, is the hypergraph formed from* $P_{\vec{\pi},2\ell}$ *by identifying the attach tuples of* $P_{\vec{\pi},2\ell}$*.*

Lemma 9. *If* π *is an (unordered) proper partition of* k *and* $\vec{\pi}$ *and* $\vec{\pi}'$ *are two orderings of* π*, then* $C_{\vec{\pi},2\ell} \cong C_{\vec{\pi}',2\ell}$*.*

Proof. Let $\vec{\pi} = (k_1, \ldots, k_t)$ and let $\eta : \{2, \ldots, t\} \to \{2, \ldots, t\}$ be any bijection of the numbers $2, \ldots, t$. We first claim that $S_{\vec{\pi}} \cong S_{(k_1, k_{\eta(2)}, \ldots, k_{\eta(t)})}$. This follows directly from the definition of the step; the bit strings can be permuted using η. That is, the isomorphism between $S_{\vec{\pi}}$ and $S_{(k_1, k_{\eta(2)}, \ldots, k_{\eta(t)})}$ is the isomorphism which takes a vertex in A in $V(S_{\vec{\pi}})$ with binary code $a_2 \ldots a_t$ to the vertex with code $a_{\eta^{-1}(2)} \ldots a_{\eta^{-1}(t)}$ in A in $S_{(k_1, k_{\eta(2)}, \ldots, k_{\eta(t)})}$ and which also takes a vertex in B_j in $S_{\vec{\pi}}$ with code $b_2 \ldots b_{j-1} b_{j+1} \ldots b_t$ to the vertex with code $b_{\eta^{-1}(2)} b_{\eta^{-1}(3)} \ldots b_{\eta^{-1}(\eta(j)-1)} b_{\eta^{-1}(\eta(j)+1)} \ldots b_{\eta^{-1}(t)}$ in $B_{\eta(j)}$ in $S_{(k_1, k_{\eta(2)}, \ldots, k_{\eta(t)})}$. To see that this bijection preserves hyperedges, let $\{\vec{a}, \vec{b_2}, \ldots, \vec{b_t}\}$ be a set of vertices in $S_{\vec{\pi}}$ where $\vec{a} \in A$ and $\vec{b_j} \in B_j$. For every j, we have $a_2 \ldots a_{j-1} a_{j+1} \ldots a_t = b_2 \ldots b_{j-1} b_{j+1} \ldots b_t$ if and only if $a_{\eta^{-1}(2)} a_{\eta^{-1}(3)} \ldots a_{\eta^{-1}(\eta(j)-1)} a_{\eta^{-1}(\eta(j)+1)} \ldots a_{\eta^{-1}(t)} = b_{\eta^{-1}(2)} b_{\eta^{-1}(3)} \ldots b_{\eta^{-1}(\eta(j)-1)} b_{\eta^{-1}(\eta(j)+1)} \ldots b_{\eta^{-1}(t)}$. This implies that $\{\vec{a}, \vec{b_2}, \ldots, \vec{b_t}\}$ is an edge of $S_{\vec{\pi}}$ if and only if the image is an edge of $S_{(k_1, k_{\eta(2)}, \ldots, k_{\eta(t)})}$.

By the previous paragraph, it suffices to prove that $C_{\vec{\pi},2\ell} \cong C_{(k_t, k_2, \ldots, k_{t-1}, k_1), 2\ell}$ to complete the proof of the lemma. Indeed, if $\vec{\pi}' = (k_{f(1)}, \ldots, k_{f(t)})$ with $f(1) > 1$, then the transformations $(1, \ldots, t) \to (1, 2, \ldots, f(1)-1, f(1)+1, \ldots, t, f(1)) \to (f(1), 2, \ldots, f(1)-1, f(1)+1, \ldots, t, 1) \to (f(1), \ldots, f(t))$ show that $C_{\vec{\pi},2\ell} \cong C_{\vec{\pi}',2\ell}$.

The fact that $C_{\vec{\pi},2\ell} \cong C_{(k_t, k_2, \ldots, k_{t-1}, k_1), 2\ell}$ is easy to see in Figure 4. In Figure 4, consider swapping the diamond and square vertices. This changes the path, but the cycle is Figure 4 with the diamond vertices on the ends identified, so swapping the diamond and square vertices preserves the cycle. In general, as discussed before, the step $S_{\vec{\pi}}$ can be drawn similar to Figure 3 *(b)* as a collection M_0 of k-partite edges between $A^{(0)}$ and B_t and a collection M_1 of k-partite edges between $A^{(1)}$ and B_t and so the path $P_{\vec{\pi},2\ell}$ can be visualized like Figure 4. Therefore the cycle $C_{\vec{\pi},2\ell}$ consists of a list of collections of k-partite edges; $E(C_{\vec{\pi},2\ell})$ is $M_{1,0} \dot\cup M_{1,1} \dot\cup M_{2,0} \dot\cup M_{2,1} \dot\cup \cdots \dot\cup M_{\ell,0} \dot\cup M_{\ell,1}$ where $M_{i,0} \cup M_{i,1}$ is a copy of $S_{\vec{\pi}}$. But $M_{i,1} \cup M_{i+1,0}$ (modulo ℓ) forms a copy of $S_{(k_t, k_2, \ldots, k_{t-1}, k_1)}$ for all i as follows. Since bit strings

in the A-part of $M_{i,1}$ have last bit one by definition, drop the last bit. After dropping these bits, for each edge E in $M_{i,1}$, the vertices in $E \cap A$ and $E \cap B_t$ have the same code. Also, for $2 \leq j \leq t-1$ the code for the vertices $E \cap B_j$ is formed by adding a one to the bit string for $E \cap A$ and then deleting the $(j-1)$th entry. But since $E \cap A$ and $E \cap B_t$ have the same code, this is the same as adding a one to the bit string for $E \cap B_t$ and then deleting the $(j-1)$th entry. Therefore we could add a one to the end of all the codes in B_t and now have half of the edges which make up $S_{(k_t,k_2,\ldots,k_{t-1},k_1)}$. A similar argument shows that $M_{i+1,0}$ forms the other half and so $M_{i,1} \cup M_{i+1,0}$ is a copy of $S_{(k_t,k_2,\ldots,k_{t-1},k_1)}$. Thus $C_{\vec{\pi},2\ell}$ is built out of ℓ copies of $S_{(k_t,k_2,\ldots,k_{t-1},k_1)}$. □

Definition. *Let π be an (unordered) proper partition of k. The* cycle of type π and length 2ℓ*, denoted $C_{\pi,2\ell}$, is $C_{\vec{\pi},2\ell}$ where $\vec{\pi}$ is any ordering of π.*

Definition. *A* walk of type $\vec{\pi}$ and length 2ℓ in a hypergraph H *is a (possibly non-injective) function $f : V(P_{\vec{\pi},2\ell}) \to V(H)$ that preserves edges. Informally, a walk is a path where the vertices are not necessarily distinct. A* circuit of type π of length 2ℓ in a hypergraph H *is a function $f : V(C_{\pi,2\ell}) \to V(H)$ that preserves edges. Informally, a circuit is a cycle where the vertices are not necessarily distinct.*

There are two alternative definitions of the cycle of length four. First, Conlon et al. [22] defined a cycle of length four for $\pi = 1 + \cdots + 1$ by an operation called *reflection*. Our definition of $C_{1+\cdots+1,4}$ is equivalent to the definition in [22]; this can be seen by noticing that the bit strings in our definition keep track of the vertex duplications which occur during reflection.

Finally, there is a concise direct definition of the cycle of type π and length four which avoids the complexity of defining steps and paths. We will not use this shorter definition in this paper, instead working with steps, paths, and walks, but we include this short definition for completeness. Let $D_1, \ldots, D_t$ be disjoint sets of size 2^{t-1} whose elements are labeled by $(t-1)$-length binary strings. The vertex set is $D_1 \dot\cup \ldots \dot\cup D_t$. For $d_1 \in D_1, \ldots, d_t \in D_t$, make $\{d_1, \ldots, d_t\}$ a hyperedge if there exists a binary string s of length t such that the code for d_i equals the code formed by deleting the ith bit of s. The cycle for general π is formed by enlarging this cycle appropriately. Figure 5 shows cycles drawn using this definition.

3 Hypergraph Eigenvalues

This section contains the definition of the largest and second largest eigenvalues of a hypergraph with respect to π and also contains some discussion and basic facts about them.

There have been three independently developed approaches to hypergraph eigenvalues: a definition by Chung [14] and Lu and Peng [50, 51] using matrices, an approach of Friedman and Wigderson [30, 32] and Cooper and Dutle [23], and lastly the eigenvalues of the shadow graph [9, 29, 52, 53, 69]. The definitions of Friedman and Wigderson [30, 32] are most suitable for our purposes and we will use their definitions as our starting point.

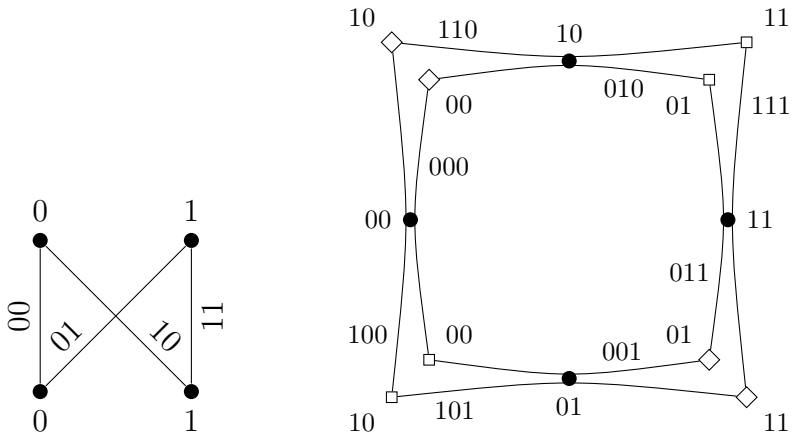


Figure 5: Alternate definition of the cycle of length four

Definition. *Let $V_1, \ldots, V_k$ be finite-dimensional vector spaces over $\mathbb{R}$. A* k-linear map *is a function $\phi : V_1 \times \cdots \times V_k \to \mathbb{R}$ such that for each $1 \leq i \leq k$, ϕ is linear in the ith coordinate. That is, for every fixed $x_i \in V_i$, $\phi(x_1, \ldots, x_{i-1}, \cdot, x_{i+1}, \ldots, x_n)$ is a linear map from V_i to $\mathbb{R}$. A k-linear map $\phi : V^k \to \mathbb{R}$ is* symmetric *if for all permutations η of $[k]$ and all $x_1, \ldots, x_k \in V$, $\phi(x_1, \ldots, x_k) = \phi(x_{\eta(1)}, \ldots, x_{\eta(k)})$.*

Definition. *Let $V_1, \ldots, V_k$ be finite-dimensional vector spaces over $\mathbb{R}$, let $B_i = \{b_{i,1}, \ldots, b_{i,\dim(V_i)}\}$ be an orthonormal basis of V_i, and let $\phi : B_1 \times \cdots \times B_k \to \mathbb{R}$ be any map.* Extending ϕ linearly to $V_1 \times \cdots \times V_k$ *means that ϕ is extended to a map $V_1 \times \cdots \times V_k \to \mathbb{R}$ where for $x_1 \in V_1, \ldots, x_k \in V_k$,*

$$\phi(x_1, \ldots, x_k) = \sum_{j_1=1}^{\dim(V_1)} \cdots \sum_{j_k=1}^{\dim(V_k)} \langle x_1, b_{1,j_1} \rangle \cdots \langle x_k, b_{k,j_k} \rangle \, \phi(b_{1,j_1}, \ldots, b_{k,j_k}). \tag{1}$$

Note that extending ϕ in this way produces a k-linear map.

Definition. ***(Friedman and Wigderson [30, 32])*** *Let H be a k-uniform hypergraph. The* adjacency map of H *is the symmetric k-linear map $\tau_H : W^k \to \mathbb{R}$ defined as follows, where W is the vector space over $\mathbb{R}$ of dimension $|V(H)|$. First, let*

$$\tau_H(e_{v_1}, \ldots, e_{v_k}) = \begin{cases} 1 & \{v_1, \ldots, v_k\} \in E(H), \\ 0 & \text{otherwise}, \end{cases} \qquad \text{for all } v_1, \ldots, v_k \in V(H),$$

where e_v denotes the indicator vector of the vertex v, that is the vector which has a one in coordinate v and zero in all other coordinates. We have defined the value of τ_H when the inputs are standard basis vectors of W. Extend τ_H to all the domain linearly as in (1)*, where for all i, $\{b_{i,1}, \ldots, b_{i,\dim(W)}\} = \{e_1, \ldots, e_{\dim(W)}\}$ is the standard basis of W.*

Definition. *Let $W_1, \ldots, W_k$ be finite dimensional vector spaces over $\mathbb{R}$ and let $\phi : W_1 \times \cdots \times W_k \to \mathbb{R}$ be a k-linear map. The* spectral norm of ϕ *is*

$$\|\phi\| = \sup_{\substack{x_i \in W_i \\ \|x_i\|=1}} |\phi(x_1, \ldots, x_k)| .$$

Before defining the first and second largest eigenvalue of H with respect to a general partition π, we give the definitions when $\pi = 1+\cdots+1$, that is π is the partition into k ones. For codegree regular hypergraphs with loops (see the definition of loops in the discussion that follows), these are exactly the definitions of Friedman and Wigderson [30, 32].

Definition. *Let H be an n-vertex, k-uniform hypergraph, let W be the vector space over $\mathbb{R}$ of dimension n, and let $J : W^k \to \mathbb{R}$ be the all-ones map. That is, if $e_{i_1}, \ldots, e_{i_k}$ are any standard basis vectors of W, then $J(e_{i_1}, \ldots, e_{i_k}) = 1$, and J is extended linearly to all of the domain as in* (1).

The largest eigenvalue of H with respect to $\pi = 1+\cdots+1$ *is* $\|\tau_H\|$ *and the* second largest eigenvalue of H with respect to $\pi = 1 + \cdots + 1$ *is* $\left\|\tau_H - \frac{k!|E(H)|}{n^k} J\right\|$.

For general $\vec{\pi} = (k_1, \ldots, k_t)$ we define a variant $\tau_{\vec{\pi}}$ of the adjacency map τ which is a t-linear map with vector space dimensions given by k_i. For example, when $\vec{\pi} = (2, 1, 1)$ the map $\tau_{\vec{\pi}}$ is a map $V \times W \times W \to \mathbb{R}$ where $\dim(W) = n$ (and we think of it as indexed by vertices) and $\dim(V) = n^2$ (and we think of it as indexed by pairs of vertices). In general, we will be working with vector spaces of dimension n^{k_i} and indexed by k_i-tuples of vertices. The language of tensors and the tensor product elegantly captures this arrangement of vector spaces. There are many possible equivalent definitions of tensors and tensor products, but for finite dimensional vector spaces the tensor product is particularly simple. We will not be using advanced facts about tensors; the next definition contains everything we will need.

Definition. *Let V and W be finite dimensional vector spaces over $\mathbb{R}$ of dimension n and m respectively. The* tensor product of V and W, *written $V \otimes W$, is the vector space over $\mathbb{R}$ of dimension nm.*

Tensor Product Facts. By definition, a tensor in $V \otimes W$ is a vector with nm coordinates. Rather than talk about the (i, j)th coordinate of this vector, we use the $\otimes$ symbol to label certain tensors in $V \otimes W$. For $1 \le i \le n$ and $1 \le j \le m$, denote by $e_i \otimes e'_j$ the tensor in $V \otimes W$ which has a one in the (i, j)th coordinate and a zero everywhere else. Therefore, the collection of tensors (vectors) $\{e_i \otimes e'_j : 1 \le i \le n, 1 \le j \le m\}$ forms an orthonormal basis of the vector space $V \otimes W$. More generally, if $v \in V$ and $w \in W$, denote by $v \otimes w$ the vector of dimension nm which is $\sum_{i,j} \langle v, e_i \rangle \langle w, e'_j \rangle \; e_i \otimes e'_j$. In other words, $v \otimes w$ is the vector of dimension nm where the (i, j)th coordinate equals the ith coordinate of v times the jth coordinate of w. Not all tensors in $V \otimes W$ can be written as $v \otimes w$ for some vectors v and w; for example, $e_1 \otimes e'_1 + e_1 \otimes e'_2 + e_2 \otimes e'_2$ is not equal to $v \otimes w$ for any v or w. A tensor which can be expressed as $v \otimes w$ for some v and w is called *simple* or *rank 1*.

A typical tensor a in $V \otimes W$ has the form $a = \sum_{i=1}^{n} \sum_{j=1}^{m} \alpha_{i,j} (e_i \otimes e'_j)$, where $\alpha_{i,j} \in \mathbb{R}$. That is, a tensor is a linear combination of the basis vectors, which have been labeled by

$e_i \otimes e'_j$. The length of a tensor is the length of the vector in the vector space $V \otimes W$. Thus the length of a is $\left(\sum_{i=1}^{n} \sum_{j=1}^{m} \alpha_{i,j}^2\right)^{1/2}$.

Let k be a positive integer. For $k \geq 3$, the tensor product $V^{\otimes k}$ is defined as $V \otimes V^{\otimes(k-1)}$ with $V^{\otimes 2} = V \otimes V$. In other words, the tensor product $V^{\otimes k}$ is the (unique) vector space over $\mathbb{R}$ of dimension n^k, where standard basis vectors are labeled by expressions $e_{i_1} \otimes \cdots \otimes e_{i_k}$.

We are now ready to define the map $\tau_{\vec{\pi}}$ and then the first and second largest eigenvalue of H with respect to π for a general π. In the definition, think of the tensor product $W^{\otimes k_i}$ as a vector space of dimension $\dim(W)^{k_i}$ indexed by ordered k_i-sets of vertices.

Definition. *Let W be a finite dimensional vector space over $\mathbb{R}$, let $\sigma : W^k \to \mathbb{R}$ be any k-linear function, and let $\vec{\pi}$ be a proper ordered partition of k, so $\vec{\pi} = (k_1, \ldots, k_t)$ for some integers $k_1, \ldots, k_t$ with $t \geq 2$. Now define a t-linear function $\sigma_{\vec{\pi}} : W^{\otimes k_1} \times \cdots \times W^{\otimes k_t} \to \mathbb{R}$ by first defining $\sigma_{\vec{\pi}}$ when the inputs are basis vectors of $W^{\otimes k_i}$ and then extending linearly as in* (1)*. For each i, $B_i = \{b_{i,1} \otimes \cdots \otimes b_{i,k_i} : b_{i,j}$ is a standard basis vector of $W\}$ is a basis of $W^{\otimes k_i}$, so for each i, pick $b_{i,1} \otimes \cdots \otimes b_{i,k_i} \in B_i$ and define*

$$\sigma_{\vec{\pi}}\left(b_{1,1} \otimes \cdots \otimes b_{1,k_1}, \ldots, b_{t,1} \otimes \cdots \otimes b_{t,k_t}\right) = \sigma(b_{1,1}, \ldots, b_{1,k_1}, \ldots, b_{t,1}, \ldots, b_{t,k_t}).$$

Now extend $\sigma_{\vec{\pi}}$ linearly to all of the domain as in (1)*. $\sigma_{\vec{\pi}}$ will be t-linear since σ is k-linear.*

Definition. *Let H be a k-uniform hypergraph and let $\tau = \tau_H$ be the (k-linear) adjacency map of H. Let π be any (unordered) partition of k and let $\vec{\pi}$ be any ordering of π. The* largest and second largest eigenvalues of H with respect to π*, denoted $\lambda_{1,\pi}(H)$ and $\lambda_{2,\pi}(H)$, are defined as*

$$\lambda_{1,\pi}(H) := \|\tau_{\vec{\pi}}\| \qquad \text{and} \qquad \lambda_{2,\pi}(H) := \left\|\tau_{\vec{\pi}} - \frac{k!|E(H)|}{n^k} J_{\vec{\pi}}\right\|.$$

Both $\lambda_{1,\pi}(H)$ and $\lambda_{2,\pi}(H)$ are well defined since τ and J are symmetric maps.

Remarks.

- Bilinear maps $\phi : W^2 \to \mathbb{R}$ correspond exactly to square matrices. Indeed, given a square matrix A, $\phi(x, y) = x^T A y$ is a bilinear map. In the other direction, for any bilinear map $\phi : W^2 \to \mathbb{R}$, there exists some square matrix A such that $\phi(x, y) = x^T A y$. The entries of A are the values $\phi(e_i, e_j)$ where e_i is the vector with one in the ith coordinate and zero everywhere else. Similarly, a k-linear map $\phi : W^k \to \mathbb{R}$ is determined by the values it takes when the inputs range over a basis of W so ϕ corresponds to a k-dimensional array of real numbers. These k-dimensional arrays are sometimes called *hypermatrices*.

- The previous paragraph combined with the definition of adjacency map shows that the adjacency map of a graph is identical to the adjacency matrix, which is a two-dimensional array of zero/ones with rows and columns indexed by vertices. For general

k, the adjacency map of a k-uniform hypergraph corresponds to a k-dimensional array of zero/ones with each axis labeled by the vertices. A location in this k-dimensional array is one if the vertices of its coordinates form a hyperedge, otherwise it is zero.

- When $k = 2$ and $\pi = 1 + 1$, the spectral norm of τ_H of a graph H exactly corresponds to the largest eigenvalue of H in absolute value in the usual definition of eigenvalues of graphs. The symmetric bilinear map τ_H is the map $\tau_H(x, y) = x^T Ay$ where A is the adjacency matrix of H. From linear algebra, we know that the spectral norm of A is the largest eigenvalue in absolute value.

- When $k = 2$ and $\pi = 1 + 1$, our definition of second largest eigenvalue of a d-regular graph H is equivalent to the usual definition of second largest eigenvalue in absolute value of the adjacency matrix A. If G is a d-regular graph, then $2|E(H)|/n^2 = \frac{d}{n}$, so our definition of the second largest eigenvalue of G is the spectral norm of $\tau_G - \frac{d}{n}J$. This bilinear map corresponds to the matrix $A - \frac{d}{n}J$ where A is the adjacency matrix and J is now the all-ones matrix. The largest eigenvalue of $A - \frac{d}{n}J$ in absolute value is the second largest eigenvalue of A in absolute value, and this equals the spectral norm of the respective map.

- The previous paragraph extends to d-codegree regular hypergraphs with loops, which were studied by Friedman and Wigderson [30, 32]. A *k-uniform hypergraph with loops* is a pair of sets $(V(H), E(H))$, where $V(H)$ is any finite set and $E(H)$ is a set of k-multisets of elements from $V(H)$. Informally, multiple edges are not allowed, each edge has size exactly k, and a vertex can be repeated multiple times inside an edge. A k-uniform hypergraph H with loops is *d-codegree regular* if for every $(k-1)$-multiset S of $V(H)$, there are exactly d hyperedges containing S. For such a hypergraph, $k!|E(H)|/n^k = d/n$ so $\lambda_{2,1+\cdots+1}(H) = \left\|\tau_H - \frac{d}{n}J\right\|$. This is exactly the definition of Friedman and Wigderson [30, 32]. In addition, they proved that if H is d-codegree regular with loops, then $\lambda_{1,1+\cdots+1}(H) = dn^{(k-2)/2}$ and is achieved by the all-ones vectors scaled to unit length. This generalizes the well known fact that the largest eigenvalue of a d-regular graph is d with eigenvector the all-ones vector. Friedman and Wigderson [30, 32] also proved several facts about $\lambda_{2,1+\cdots+1}(H)$ including upper and lower bounds, an Expander Mixing Lemma which we generalize to all π in Theorem 10, and the value of $\lambda_{2,1+\cdots+1}$ on the random hypergraph.

- If π' is a refinement of π, then it is easy to see that $\lambda_{1,\pi'}(H) \leq \lambda_{1,\pi}(H)$ and $\lambda_{2,\pi'}(H) \leq \lambda_{2,\pi}(H)$ for every H. For example, let $\vec{\pi}' = (2, 2, 1)$ and $\vec{\pi} = (3, 2)$. Take $x \in W \otimes W$, $y \in W \otimes W$, and $z \in W$ so that $\tau_{\vec{\pi}'}(x, y, z) = \lambda_{1,\pi'}(H)$, that is x, y, z achieve the supremum in the definition of spectral norm. But now $\tau_{\vec{\pi}}(x \otimes z, y) = \tau_{\vec{\pi}'}(x, y, z)$ so the spectral norm of $\tau_{\vec{\pi}}$ is either equal to $\tau_{\vec{\pi}'}(x, y, z)$ or larger ($x \otimes z$ is a unit length tensor since both x and z are unit length). This generalizes to all partition refinements and also generalizes to the map $\sigma = \tau - \frac{k!|E(H)|}{n^k}J$. Thus refining the partition does not increase $\lambda_{1,\pi}$ or $\lambda_{2,\pi}$, which matches exactly with the arrangement of quasirandom properties in Figure 1: $\texttt{Eig}[\pi] \Rightarrow \texttt{Eig}[\pi']$.

4 $\texttt{Eig}[\pi] \Rightarrow \texttt{Expand}[\pi]$

In this section we prove a generalization of the graph Expander Mixing Lemma which relates spectral and expansion properties of graphs. It was first discovered independently by Alon and Milman [4] and Tanner [72]. For background on graph expansion and eigenvalues, see [1, 7, 38]. The following theorem extends on the hypergraph Expander Mixing Lemma of Friedman and Widgerson [30, 32], which applied for $\pi = 1 + \cdots + 1$. The theorem is stated for ordered partitions $\vec{\pi}$, but trivially gives the same result for any ordering $\vec{\pi}$ of a partition π.

Theorem 10. *(Hypergraph Expander Mixing Lemma) Let H be an n-vertex, k-uniform hypergraph. Let $\vec{\pi} = (k_1, \ldots, k_t)$ be a proper ordered partition of k and let $S_i \subseteq \binom{V(H)}{k_i}$ for $1 \leq i \leq t$. Then*

$$\left| e(S_1, \ldots, S_t) - \frac{k!|E(H)|}{n^k} |S_1| \cdots |S_t| \right| \leq \lambda_{2,\pi}(H) \sqrt{|S_1| \cdots |S_t|},$$

where $e(S_1, \ldots, S_t)$ is the number of ordered tuples $(s_1, \ldots, s_t)$ such that $s_1 \cup \cdots \cup s_t \in E(H)$ and $s_i \in S_i$.

Proof. Let $q = \frac{k!|E(H)|}{n^k}$, let τ_H be the adjacency map of H, and let $\sigma = \tau_H - qJ$. It is easy to see that by definition, $(\tau - qJ)_{\vec{\pi}} = \tau_{\vec{\pi}} - qJ_{\vec{\pi}}$, so $\lambda_{2,\pi}(H) = \|\sigma_{\vec{\pi}}\|$. Let $\chi_{S_i} \in W^{k_i \otimes}$ be the indicator tensor of S_i. If we let $V(H) = [n]$, then

$$\chi_{S_i} = \sum_{\substack{\{s_1, \ldots, s_{k_i}\} \in S_i \\ s_1 \leq \cdots \leq s_{k_i}}} (e_{s_1} \otimes \cdots \otimes e_{s_{k_i}}).$$

By the linearity of $\sigma_{\vec{\pi}}$ and the definition of $J_{\vec{\pi}}$,

$$\sigma_{\vec{\pi}}(\chi_{S_1}, \ldots, \chi_{S_t}) = \tau_{\vec{\pi}}(\chi_{S_1}, \ldots, \chi_{S_t}) - qJ_{\vec{\pi}}(\chi_{S_1}, \ldots, \chi_{S_t}) = e(S_1, \ldots, S_t) - q \prod_{i=1}^{t} |S_i|.$$

Before upper bounding this by $\lambda_{2,\pi}(H)$, we must scale each indicator tensor to be unit length. Since $\{e_{j_1} \otimes \cdots \otimes e_{j_{k_i}} : 1 \leq j_1, \ldots, j_{k_i} \leq n\}$ forms a basis of $W^{\otimes k_i}$, we have $\|\chi_{S_i}\| = \sqrt{|S_i|}$. Thus

$$\left| \sigma_{\vec{\pi}} \left(\frac{\chi_{S_1}}{\|\chi_{S_1}\|}, \ldots, \frac{\chi_{S_t}}{\|\chi_{S_t}\|} \right) \right| \leq \|\sigma_{\vec{\pi}}\| = \lambda_{2,\pi}(H).$$

Consequently,

$$|\sigma_{\vec{\pi}}(\chi_{S_1}, \ldots, \chi_{S_t})| \leq \lambda_{2,\pi}(H) \, \|\chi_{S_1}\| \cdots \|\chi_{S_t}\| = \lambda_{2,\pi}(H) \sqrt{|S_1| \cdots |S_t|},$$

and the proof is complete. □

Note that we could get an even tighter bound in Theorem 10 by projecting each χ_{S_i} onto the space perpendicular to the all-ones vector to keep the value of $\sigma_{\vec{\pi}}$ the same but shorten the length of the vectors given as input. Friedman and Widgerson [30, 32] compute this tighter bound for $\pi = 1 + \cdots + 1$; we leave this as an exercise to the reader for general π because the tighter bound is not required to prove $\texttt{Eig}[\pi] \Rightarrow \texttt{Expand}[\pi]$.

Lemma 11. *Let $\mathcal{H} = \{H_n\}$ be a sequence of k-uniform hypergraphs with $|V(H_n)| = n$ and $|E(H_n)| \geq p\binom{n}{k} + o(n^k)$. Let τ_n be the adjacency map of H_n and let $\vec{\pi} = (k_1, \ldots, k_t)$ be a proper ordered partition of k. If $\lambda_{1,\pi}(H_n) = pn^{k/2} + o(n^{k/2})$, then $|E(H_n)| = p\binom{n}{k} + o(n^k)$.*

Proof. Throughout this proof the subscripts on n are dropped for simplicity. Let W be the vector space over $\mathbb{R}$ of dimension n. For $1 \leq i \leq t$, let $\vec{1}_{k_i}$ denote the all-ones vector in $W^{\otimes k_i}$, so $\|\vec{1}_{k_i}\| = n^{k_i/2}$. Then

$$\begin{aligned}\tau_{\vec{\pi}}\left(\frac{\vec{1}_{k_1}}{n^{k_1/2}}, \ldots, \frac{\vec{1}_{k_t}}{n^{k_t/2}}\right) &= \frac{1}{n^{k/2}}\tau_{\vec{\pi}}\left(\vec{1}_{k_1}, \ldots, \vec{1}_{k_t}\right) = \frac{1}{n^{k/2}}\tau(\vec{1}_1, \ldots, \vec{1}_1) \\ &= \frac{1}{n^{k/2}} \sum_{i_1,\ldots,i_k=1}^{n} \tau(e_{i_1}, \ldots, e_{i_k}) \\ &= \frac{1}{n^{k/2}} k!|E(H)|.\end{aligned}$$

Thus the spectral norm of $\tau_{\vec{\pi}}$ is at least $k!|E(H)|/n^{k/2}$, so

$$pn^{k/2} \leq \frac{k!|E(H)|}{n^{k/2}} + o(n^{k/2}) \leq \|\tau_{\vec{\pi}}\| + o(n^{k/2}) = pn^{k/2} + o(n^{k/2})$$

This implies equality (up to $o(n^{k/2})$) throughout the above expression. In particular, $|E(H_n)| = p\binom{n}{k} + o(n^k)$. □

Proof that $\texttt{Eig}[\pi] \Rightarrow \texttt{Expand}[\pi]$. First, $\texttt{Eig}[\pi]$ contains the assertion that $\lambda_{1,\pi}(H_n) = pn^{k/2} + o(n^{k/2})$ which by Lemma 11 implies $|E(H_n)| = p\binom{n}{k} + o(n^k)$. Consequently, $k!|E(H_n)|/n^k = (1 + o(1))p$ and Theorem 10 imply that

$$\left|e(S_1, \ldots, S_t) - (1 + o(1))p|S_1| \cdots |S_t|\right| \leq \lambda_{2,\pi}(H)\sqrt{|S_1| \cdots |S_t|} \tag{2}$$

for any choice of $S_i \subseteq \binom{V(H_n)}{k_i}$, $i = 1, \ldots, t$. Since π is a partition of k, $\sqrt{|S_1| \cdots |S_t|} = O(n^{k/2})$. Also, $\texttt{Eig}[\pi]$ states that $\lambda_{2,\pi}(H) = o(n^{k/2})$. Thus (2) becomes

$$\left|e(S_1, \ldots, S_t) - p|S_1| \cdots |S_t|\right| = o(n^k),$$

which proves $\texttt{Expand}[\pi]$. □

5 Expand$[\pi] \Rightarrow$ Count[π-linear]

In this section we prove Expand$[\pi] \Rightarrow$ Count[π-linear] by proving an embedding lemma for hypergraphs. The proof below is a generalization of an argument by Kohayakawa et al. [41] who proved it in the special case of linear hypergraphs. The lemma below is proved for ordered partitions $\vec{\pi}$, but it is easy to see that the lemma is independent of the ordering chosen for $\vec{\pi}$.

Proposition 12. *Let $\vec{\pi} = (k_1, \ldots, k_t)$ be a proper ordered partition of k, let $0 < p < 1$, and let F be any fixed k-uniform, π-linear hypergraph with f vertices and m edges.*

Let $\mathcal{H} = \{H_n\}_{n\to\infty}$ be a sequence of k-uniform hypergraphs with $|V(H_n)| = n$, $|E(H_n)| = p\binom{n}{k} + o(n^k)$, and for which Expand$[\vec{\pi}]$ holds. That is, for every $S_1 \subseteq \binom{V(H)}{k_1}, \ldots, S_t \subseteq \binom{V(H)}{k_t}$, we have $e(S_1, \ldots, S_t) = p|S_1| \cdots |S_t| + o(n^k)$. Then the number of labeled copies of F in H equals $p^m n^f + o(n^f)$.

Proof. The proof is by induction on the number of edges of F. If F has zero or one edge, then the result is trivial. So assume that F has at least two edges and let E be the last edge in the ordering provided by the π-linearity of F. Let F_* be the hypergraph formed by deleting all vertices of E from F. Let Q_* be a labeled copy of F_* in H by which we mean that Q_* is an injective edge preserving map $Q_* : V(F_*) \to V(H)$. We can count the number of labeled copies of F in H by counting for each Q_*, the number of ways Q_* extends to a labeled copy of F. More precisely, we count the number of edge preserving injections $Q : V(F) \to V(H)$ which when restricted to $V(F_*)$ match the injection Q_*. Since F is π-linear, the edge E can be divided into $A_1, \ldots, A_t$ such that $|A_i| = k_i$ and the edges of F intersecting E can be divided into sets $R_1, \ldots, R_t$ such that every edge in R_i intersects E in a subset of A_i.

Consider some Q_* in H; we will count how many ways it extends to a labeled copy of F. For $1 \le i \le t$, define $S_i(Q_*)$ to be the following collection of k_i-sets. Let $Y \subseteq V(H)$ be a set of k_i vertices and add Y to $S_i(Q_*)$ if $Y \cap Im(Q_*) = \emptyset$ and there exists an edge preserving injection $V(F_*) \cup A_i \to Im(Q_*) \cup Y$ which when restricted to $V(F_*)$ matches the map Q_*. More informally, $S_i(Q_*)$ consists of all k_i-sets Y of vertices which can be used to extend Q_* to embed a labeled copy of $F_* \cup R_i$.

Every edge counted by $e(S_1(Q_*), \ldots, S_t(Q_*))$ creates several labeled copies of F which extend Q_*. First, let Δ_i be the number of edge preserving bijections $V(F_*) \cup A_i \to V(F_*) \cup A_i$ which are the identity map when restricted to $V(F_*)$. More informally, if we are given a non-labeled $F_* \cup R_i$ together with a labeling of the vertices of F_*, Δ_i is the number of ways of labeling the vertices of A_i. The numbers $\Delta_1, \ldots, \Delta_t$ are fixed numbers depending only on F; Δ_i depends on the way that edges in R_i intersect. For example, if every edge in R_i meets E in exactly A_i, then $\Delta_i = k_i!$. If the edges in R_i intersect differently, Δ_i will change but still depend only on F. Now we count the number of edge preserving injections $Q : V(F) \to V(H)$ where $Q|_{V(F_*)} = Q_*$ as follows. First pick an edge to use for E; this consists of picking one of the edges which take a k_1-set Y_1 from $S_1(Q_*)$, a k_2-set Y_2 from $S_2(Q_*)$, and so on. There are exactly $e(S_1(Q_*), \ldots, S_t(Q_*))$ such edges. We are embedding a labeled copy of F so next we order the vertices inside the set Y_1 chosen from $S_1(Q_*)$; there are Δ_1 ways of ordering the vertices of Y_1. Similarly, we order the vertices inside the other sets chosen from $S_i(Q_*)$

for a total of $\prod \Delta_i$ orderings. These are all the labeled copies of F which extend Q_*, so there are exactly $e(S_1(Q_*),\ldots,S_t(Q_*))\prod \Delta_i$ labeled copes of F extending Q_*, i.e. exactly $e(S_1(Q_*),\ldots,S_t(Q_*))\prod \Delta_i$ edge preserving injections $V(F)\to V(H)$ which when restricted to $V(F_*)$ are Q_*. By assupmtion,

$$e\big(S_1(Q_*),\ldots,S_t(Q_*)\big) = p|S_1(Q_*)|\cdots|S_t(Q_*)| + o(n^k).$$

Therefore, we can count the number of labeled copies of F in H by summing the value of $e(S_1(Q_*),\ldots,S_t(Q_*))\prod \Delta_i$ over all Q_* in H. By the notation $\#\{F \text{ in } H\}$ we mean the number of labeled copies of F in H.

$$\begin{aligned}
\#\{F \text{ in } H\} &= \sum_{Q_*} e(S_1(Q_*),\ldots,S_t(Q_*))\prod_i \Delta_i \\
&= \sum_{Q_*}\left(p\,|S_1(Q_*)|\cdots|S_t(Q_*)|\prod_i \Delta_i + o(n^k)\right) \\
&= p\sum_{Q_*}|S_1(Q_*)|\cdots|S_t(Q_*)|\prod_i \Delta_i + o(n^f), \qquad (3)
\end{aligned}$$

since the number of Q_* is bounded above by $n^{|V(F_*)|} = n^{f-k}$ so summing the term $o(n^k)$ over Q_* is bounded by $o(n^f)$. Let F_- be the hypergraph formed by removing the edge E from F but keeping the same vertex set. Then similarly to the above, the number of labeled copies of F_- extending Q_* is $|S_1(Q_*)|\cdots|S_t(Q_*)|\prod \Delta_i$ since every labeled copy of F_- is formed by picking a set Y_1 from $S_1(Q_*)$, ordering it in Δ_i ways, picking a set Y_2 from $S_2(Q_*)$ and ordering it, and so on. Therefore the number of labeled copies of F_- in H is counted by summing over Q_* and counting $|S_1(Q_*)|\cdots|S_t(Q_*)|\prod \Delta_i$. Thus (3) continues as

$$\#\{F \text{ in } H\} = p\,\#\{F_- \text{ in } H\} + o(n^f). \qquad (4)$$

By hypothisis, F is π-linear so F_- is π-linear. Thus by induction the number of labeled copies of F_- in H is $p^{m-1}n^f + o(n^f)$. Inserting this into (4) shows that the number of labeled copies of F in H is $p^m n^f + o(n^f)$. □

6 $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$

In Section 6.1, we state two crucial propositions on products and powers of multilinear maps and then show that these two propositions imply that $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$. The remainder of Section 6 is concerned with the proof of one of these propositions (the other is proved in the appendix). In Section 6.2 we discuss the special case of $\vec{\pi} = (3,2,1)$, in Section 6.3 we state the definitions of products and powers of multilinear maps, and finally in Section 6.4 we show that the powers of the adjacency map counts walks in a hypergraph.

6.1 Overview

First, let us recall the proof of $\texttt{Cycle}_4[1,1] \Rightarrow \texttt{Eig}[1,1]$ for graphs. If A is the adjacency matrix of a graph G, then A^4 counts walks of length 4 in the sense that the (i,j)-th entry of A^4 is the number of walks of length 4 between i and j. The trace of A^4 is then the number of circuits of length 4 in G. The trace of a square real symmetric matrix is the sum of its eigenvalues. If G is d-regular, then the largest eigenvalue of A^4 is d^4 so if the number of circuits of length four in G is $d^4 + o(n^4)$, then the trace of A^4 is $d^4 + o(n^4)$ which implies that all the eigenvalues of A besides d are $o(n)$. The proof for non-regular graphs is a straightforward generalization of this. Our proof for hypergraphs follows the same outline once some algebraic facts about multilinear maps are proved.

The first step in the general proof of $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$ is to define the product of two multilinear maps. This is a simple (although slightly verbose) definition which is an easy generalization of matrix multiplication. Next, we define the power of a multilinear map and show that if τ is the adjacency map of a hypergraph, then powers of $\tau_{\vec{\pi}}$ count walks. Since circuits are walks where the end points are identified, counting circuits will correspond to the trace of a power of $\tau_{\vec{\pi}}$. In Section 6.3, we define for any t-linear map ψ a square symmetric matrix $A[\psi^{2^{t-1}}]$. We write $A[\psi^{2^{t-1}}]^\ell$ for the ℓ-th matrix power of $A[\psi^{2^{t-1}}]$. The key proposition is the following.

Proposition 13. *Let H be a k-uniform hypergraph, let $\vec{\pi}$ be a proper ordered partition of k, and let $\ell \geq 2$ be an integer. Let $\tau = \tau_H$ be the adjacency map of H. Then $Tr\left[A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell\right]$ is the number of labeled circuits of type $\vec{\pi}$ and length 2ℓ in H.*

In Section 6.2, we sketch a proof of this proposition for the special case of $\vec{\pi} = (3,2,1)$ and in Section 6.4 give the proof for all $\vec{\pi}$. Next we need an algebraic property of the eigenvalues of multilinear maps. As part of their proof for graphs, Chung, Graham, and Wilson [18] proved that in a graph sequence satisfying $\texttt{Eig}$, the distance between the all-ones vector and the eigenvector corresponding to the largest eigenvalue is $o(1)$ (see the bottom of page 350 in [18]). The reason for this is that if A is the adjacency matrix and v the unit length eigenvector corresponding to the largest eigenvalue, then the second largest eigenvalue of A is the spectral norm of $A - \lambda_1 vv^T$. But as we saw in Section 4, proving $\texttt{Expand}$ requires bounding the spectral norm of $A - pn\hat{1}\hat{1}^T$, where $\hat{1}$ is the all-ones vector scaled to unit length (note that $n\hat{1}\hat{1}^T = J$). Chung, Graham, and Wilson [18] proved that $\left\|v - \hat{1}\right\| = o(1)$ to conclude that $A - \lambda_1 vv^T$ and $A - pn\hat{1}\hat{1}^T$ are almost the same map so their spectral norms are asymptotically equal. The following proposition generalizes this to hypergraphs.

Proposition 14. *Let $\{\psi_r\}_{r\to\infty}$ be a sequence of symmetric k-linear maps, where $\psi_r : V_r^k \to \mathbb{R}$, V_r is a vector space over $\mathbb{R}$ of finite dimension, and $\dim(V_r) \to \infty$ as $r \to \infty$. Let $\hat{1}$ denote the all-ones vector in V_r scaled to unit length and let $J : V_r^k \to \mathbb{R}$ be the k-linear all-ones map. Let π be a proper (unordered) partition of k, and assume that for every ordering*

$\vec{\pi}$ of π,

$$\lambda_1(A[\psi_{\vec{\pi}}^{2^{t-1}}]) = (1+o(1))\psi\left(\hat{1},\ldots,\hat{1}\right)^{2^{t-1}},$$
$$\lambda_2(A[\psi_{\vec{\pi}}^{2^{t-1}}]) = o\left(\lambda_1(A[\psi_{\vec{\pi}}^{2^{t-1}}])\right).$$

Then for every ordering $\vec{\pi}$ of π,

$$\|\psi_{\vec{\pi}} - qJ_{\vec{\pi}}\| = o(\psi(\hat{1},\ldots,\hat{1})),$$

where $q = \dim(V_r)^{-k/2}\psi(\hat{1},\ldots,\hat{1})$.

For graphs, we will show later that $A[\tau^2]$ is the adjacency matrix squared, so Proposition 14 states that the spectral norm of $A - \frac{2|E(G)|}{n^2}J$ is little-o of the square root of the largest eigenvalue of A^2, exactly the fact proved by Chung, Graham, and Wilson [18]. Note that in graphs, if the graph sequence is assumed to be d-regular then Proposition 14 becomes trivial since the spectral norm of $A - \frac{d}{n}J$ equals the second largest eigenvalue of A. Similarly, if Theorem 4 is restricted to hypergraph sequences which are codegree regular, the proof of Proposition 14 becomes much simpler. Finally, we need one last simple property of multilinear maps.

Lemma 15. *Let $V_1,\ldots,V_t$ be finite dimensional vector spaces over $\mathbb{R}$ and let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map. Then $\|\phi\|^{2^{t-1}} \le \lambda_1(A[\phi^{2^{t-1}}])$.*

Propositions 13 and 14 and Lemma 15 combine to prove that $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$ for any proper partition π as follows. The proof can be understood without knowing the definition of $A[\tau_{\vec{\pi}}^{2^{t-1}}]$ (which appears in Section 6.3): $A[\tau_{\vec{\pi}}^{2^{t-1}}]$ is a square symmetric matrix.

Proof that $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$. Let $\mathcal{H} = \{H_n\}_{n\to\infty}$ be a sequence of hypergraphs and let τ_n be the adjacency map of H_n. For notational convenience, the subscript on n is dropped below. Throughout this proof, we use $\hat{1}$ to denote the all-ones vector scaled to unit length. Wherever we use the notation $\hat{1}$, it is the input to a multilinear map and so $\hat{1}$ denotes the all-ones vector in the appropriate vector space corresponding to whatever space the map is expecting as input. This means that in equations below $\hat{1}$ can stand for different vectors in the same expression, but attempting to subscript $\hat{1}$ with the vector space (for example $\hat{1}_{V_3}$) would be notationally awkward.

The proof that $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$ comes down to checking the conditions of Proposition 14. Let $\vec{\pi}$ be any ordering of the entries of π. We will show that the first and second largest eigenvalues of $A = A[\tau_{\vec{\pi}}^{2^{t-1}}]$ are separated. Let $m = |E(C_{\pi,4\ell})| = 2\ell 2^{t-1}$ and note that $|V(C_{\pi,4\ell})| = mk/2$ since $C_{\pi,4\ell}$ is two-regular. A is a square symmetric real valued matrix, so let $\mu_1,\ldots,\mu_d$ be the eigenvalues of A arranged so that $|\mu_1| \ge \cdots \ge |\mu_d|$, where $d = \dim(A)$. The eigenvalues of $A^{2\ell}$ are $\mu_1^{2\ell},\ldots,\mu_d^{2\ell}$ and the trace of $A^{2\ell}$ is $\sum_i \mu_i^{2\ell}$. Since all $\mu_i^{2\ell} \ge 0$, Proposition 13 and $\texttt{Cycle}_{4\ell}[\pi]$ implies that

$$\mu_1^{2\ell} + \mu_2^{2\ell} \le \mathrm{Tr}\left[A^{2\ell}\right] = \#\{\text{possibly degenerate } C_{\pi,4\ell} \text{ in } H_n\} \le p^m n^{mk/2} + o(n^{mk/2}). \quad (5)$$

We now verify the conditions on μ_1 and μ_2 in Proposition 14, and to do that we need to compute $\tau(\hat{1},\dots,\hat{1})$. The computations in Lemma 11 show that

$$\tau(\hat{1},\dots,\hat{1}) = \tau_{\vec{\pi}}(\hat{1},\dots,\hat{1}) = \frac{k!E(H)}{n^{k/2}}. \tag{6}$$

Using $|E(H_n)| \geq p\binom{n}{k} + o(n^k)$, Lemma 15, and $\mu_1^{2\ell} \leq p^m n^{mk/2} + o(n^{mk/2})$ from (5),

$$pn^{k/2} + o(n^{k/2}) \leq \frac{k!E(H)}{n^{k/2}} = \tau_{\vec{\pi}}(\hat{1},\dots,\hat{1}) \leq \|\tau_{\vec{\pi}}\| \leq \mu_1^{1/2^{t-1}} \leq pn^{k/2} + o(n^{k/2}). \tag{7}$$

This implies equality up to $o(n^{k/2})$ throughout the above expression, so $\tau(\hat{1},\dots,\hat{1}) = pn^{k/2} + o(n^{k/2})$, $\lambda_{1,\pi}(H_n) = \|\tau_{\vec{\pi}}\| = pn^{k/2} + o(n^{k/2})$, and $\mu_1 = p^{2^{t-1}} n^{k2^{t-2}} + o(n^{k2^{t-2}})$, so $\mu_1 = (1+o(1))\tau(\hat{1},\dots,\hat{1})^{2^{t-1}}$.

Insert $\mu_1 = p^{2^{t-1}} n^{k2^{t-2}} + o(n^{k2^{t-2}})$ into (5) to show that $\mu_2 = o(n^{k2^{t-2}})$. Therefore, the conditions of Proposition 14 are satisfied, so

$$\|\tau_{\vec{\pi}} - qJ_{\vec{\pi}}\| = o(\tau(\hat{1},\dots,\hat{1})) = o(n^{k/2}),$$

where $q = n^{-k/2}\tau(\hat{1},\dots,\hat{1})$. Using (6), $q = k!|E(H)|/n^k$. Thus $\|\tau_{\vec{\pi}} - qJ_{\vec{\pi}}\| = \lambda_{2,\pi}(H_n)$ and the proof is complete. □

The above proof can be extended to even length cycles in the case when $\vec{\pi} = (k_1, k_2)$ is a partition into two parts. For these $\vec{\pi}$, the matrix $A[\tau_{\vec{\pi}}^2]$ can be shown to be positive semidefinite since $A[\tau_{\vec{\pi}}^2]$ will equal MM^T where M is the matrix associated to the bilinear map $\tau_{\vec{\pi}}$. Since $A[\tau_{\vec{\pi}}^2]$ is positive semidefinite, each $\mu_i \geq 0$ so any power of μ_i is non-negative. For partitions into more than two parts, we don't know if the matrix $A[\tau_{\vec{\pi}}^{2^{t-1}}]$ is always positive semidefinite or not.

6.2 Counting circuits in six uniform hypergraphs

In this section, we explain the key ideas behind Proposition 13 by looking at the case of six uniform hypergraphs with $\vec{\pi} = (3,2,1)$. Throughout this section, for a finite dimensional vector space W over $\mathbb{R}$, we denote by $e_1,\dots,e_{\dim(W)}$ the standard basis vectors; e_i is a vector with a one in the ith coordinate and a zero everywhere else.

First, return to matrices and bilinear maps and let A and B be square matrices and let ϕ and ψ be the associated bilinear maps so $\phi(x,y) := x^TAy$ and $\psi(x,y) := x^TBy$. Now translate the matrix product AB into the language of bilinear maps: consider the i,jth coordinate of AB:

$$\phi\psi(e_i,e_j) = e_i^TABe_j = \sum_{p=1}^{n}(e_i^TAe_p)(e_p^TBe_j) = \sum_{p=1}^{n}\phi(e_i,e_p)\psi(e_p,e_j) \tag{8}$$

since the i,j-th entry is the dot product of the ith row of A with the jth column of B and $e_i^TAe_p$ is the pth entry of the ith row of A and $e_p^TBe_j$ is the pth entry of the jth column of

B. If both ϕ and ψ are the adjacency map, then (8) counts the number of walks of length two between i and j, since the product sums over all possible vertices p.

Now consider a six-uniform, n vertex hypergraph H and let $\tau = \tau_H$ be the adjacency map of H. Let W be a vector space over $\mathbb{R}$ of dimension $n = |V(H)|$. Consider the map $\tau_{\vec{\pi}}$ defined in Section 3:

$$\tau_{\vec{\pi}} : (W \otimes W \otimes W) \times (W \otimes W) \times W \to \mathbb{R}$$

$$\tau_{\vec{\pi}}(e_u \otimes e_v \otimes e_w, e_x \otimes e_y, e_z) = \tau(e_u, e_v, e_w, e_x, e_y, e_z) = \begin{cases} 1 & \text{if } \{u, v, w, x, y, z\} \in E(H) \\ 0 & \text{otherwise} \end{cases}$$

(extended linearly to all of the domain). Motivated by the form of matrix multiplication in equation (8), define the product of $\tau_{\vec{\pi}}$ with itself as the following map. This is a special case of the more general definition in Section 6.3. Let $a, a' \in W^{\otimes 3}$ so $a \otimes a' \in W^{\otimes 6}$ and let $b, b' \in W^{\otimes 2}$ so $b \otimes b' \in W^{\otimes 4}$. Define $\tau_{\vec{\pi}} * \tau_{\vec{\pi}}$ as follows:

$$\tau_{\vec{\pi}} * \tau_{\vec{\pi}} : W^{\otimes 6} \times W^{\otimes 4} \to \mathbb{R}$$

$$\tau_{\vec{\pi}} * \tau_{\vec{\pi}}(a \otimes a', b \otimes b') := \sum_{i=1}^{n} \tau_{\vec{\pi}}(a, b, e_i)\tau_{\vec{\pi}}(a', b', e_i), \tag{9}$$

where the definition is extended linearly to all of the domain. The product $\tau_{\vec{\pi}} * \tau_{\vec{\pi}}(a \otimes a', b \otimes b')$ will count the number of labeled copies of the hypergraph in Figure 6.

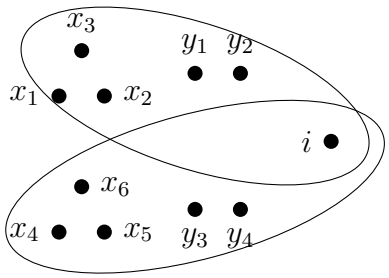


Figure 6: Hypergraph counted by $\tau_{\vec{\pi}} * \tau_{\vec{\pi}}$

Let $x_1, \ldots, x_6, y_1, \ldots, y_4$ be not necessarily distinct vertices and define $a = e_{x_1} \otimes e_{x_2} \otimes e_{x_3}$, $a' = e_{x_4} \otimes e_{x_5} \otimes e_{x_6}$, $b = e_{y_1} \otimes e_{y_2}$, and $b' = e_{y_3} \otimes e_{y_4}$. Then

$$\tau_{\vec{\pi}} * \tau_{\vec{\pi}}(a \otimes a', b \otimes b') = \sum_{i=1}^{n} \tau(e_{x_1}, e_{x_2}, e_{x_3}, e_{y_1}, e_{y_2}, e_i)\tau(e_{x_4}, e_{x_5}, e_{x_6}, e_{y_3}, e_{y_4}, e_i).$$

Thus $\tau_{\vec{\pi}} * \tau_{\vec{\pi}}(a \otimes a', b \otimes b')$ counts the number of labeled copies of Figure 6 where vertex i is allowed to range over any vertex and the other vertices are fixed.

Next, we define the power of $\tau^4_{\vec{\pi}}$ which will count steps of type $\vec{\pi}$. The following definition of $\tau^4_{\vec{\pi}}$ is a special case of a more general definition in Section 6.3. Let $a, a', a'', a''' \in W^{\otimes 3}$ and define $\tau^4_{\vec{\pi}}$ as follows:

$$\tau^4_{\vec{\pi}} : W^{\otimes 12} \to \mathbb{R}$$

$$\tau^4_{\vec{\pi}}(a \otimes a' \otimes a'' \otimes a''') := \sum_{i,j,\ell,m=1}^{n} \tau^2_{\vec{\pi}}(a \otimes a', e_i \otimes e_j \otimes e_\ell \otimes e_m)\tau^2_{\vec{\pi}}(a'' \otimes a''', e_i \otimes e_j \otimes e_\ell \otimes e_m).$$

Note that the definition still takes a similar form to (8) and (9) since the expression for $(\tau^2_{\vec{\pi}})^2$ sums over a basis of $W^{\otimes 4}$ and takes the product of $\tau^2_{\vec{\pi}}$ with $\tau^2_{\vec{\pi}}$. Expand out the definition of $\tau^2_{\vec{\pi}}$ to obtain

$$\sum_{i,j,\ell,m=1}^{n} \left(\sum_{s=1}^{n} \tau_{\vec{\pi}}(a, e_i \otimes e_j, e_s)\tau_{\vec{\pi}}(a', e_\ell \otimes e_m, e_s) \right) \tau^2_{\vec{\pi}}(a'' \otimes a''', e_i \otimes e_j \otimes e_\ell \otimes e_m)$$

$$= \sum_{i,j,\ell,m,r,s=1}^{n} \tau_{\vec{\pi}}(a, e_i \otimes e_j, e_s)\tau_{\vec{\pi}}(a', e_\ell \otimes e_m, e_s)\tau_{\vec{\pi}}(a'', e_i \otimes e_j, e_r)\tau_{\vec{\pi}}(a''', e_\ell \otimes e_m, e_r). \quad (10)$$

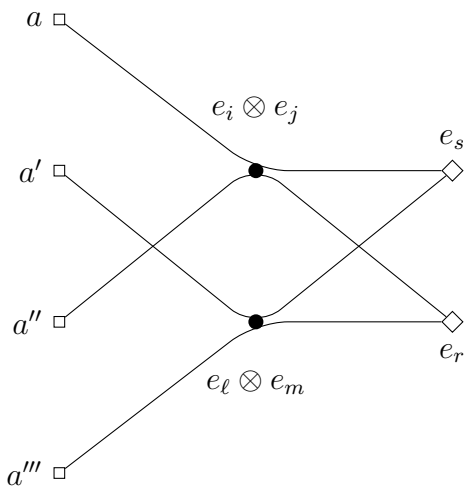


Figure 7: $S_{\vec{\pi}}$, the hypergraph counted by $\tau^4_{\vec{\pi}}$

Compare (10) and Figure 7. In the figure, each square vertex is actually a small cluster of three vertices, each filled circle is a small cluster of two vertices, and each diamond is a single vertex (recall that the step was defined by expanding Figure 3). Pick twelve vertices (four groups of three vertices) for the squares in Figure 7 and define $a, a', a'', a''' \in W^{\otimes 3}$ as indicator tensors for these vertex groups. That is, $a = e_{x_1} \otimes e_{x_2} \otimes e_{x_3}$ if x_1, x_2, x_3 are the three vertices specified for the topmost square in Figure 7. Once we specify the twelve vertices for the squares and define a, a', a'', a''', then $\tau^4_{\vec{\pi}}(a \otimes a' \otimes a'' \otimes a''')$ counts the number of copies of

the step of type $\vec{\pi}$ with the square vertices fixed and the circle and diamond vertices allowed to range over all of $V(H)$. This is easy to see from (10), since the expression sums over i, j, ℓ, m, r, s and has a one if and only if all four edges drawn in the figure appear. Also, Figure 7 is drawn slightly differently but is the same as the hypergraph in Figure 3 so is the step $S_{\vec{\pi}}$. Lastly, notice that Figure 7 is two copies of Figure 6: one copy using a, a', and s and another using a'', a''', and r. This reflects the fact that $\tau^4_{\vec{\pi}}$ sums over the product of two $\tau^2_{\vec{\pi}}$.

Next, we would like to rearrange the inputs to $\tau^4_{\vec{\pi}}$ to separate the attach tuples of the step. Note that $a \otimes a''$ is one of the attach tuples and $a' \otimes a'''$ is the other attach tuple (each attach tuple is a tuple of six vertices). We would like to define a map which is the same as the map $\tau^4_{\vec{\pi}}$, but takes the 12-vertex tuple as two 6-vertex tuples of attach tuples. Define $A[\tau^4_{\vec{\pi}}]$ to be the following matrix/bilinear map.

$$A[\tau^4_{\vec{\pi}}] : W^{\otimes 6} \times W^{\otimes 6} \to \mathbb{R}$$
$$A[\tau^4_{\vec{\pi}}](a \otimes a'', a' \otimes a''') := \tau^4_{\vec{\pi}}(a \otimes a' \otimes a'' \otimes a''').$$

With these definitions, if $w_1, w_2 \in W^{\otimes 6}$ are indicator tensors for 6-tuples of vertices, then $A[\tau^4_{\vec{\pi}}](w_1, w_2)$ counts the number of (possibly degenerate) steps of type $\vec{\pi}$ with specified attach tuples w_1 and w_2. In addition, since $A[\tau^4_{\vec{\pi}}]$ is a bilinear map $W^{\otimes 6} \times W^{\otimes 6} \to \mathbb{R}$, it corresponds to a square matrix. Also, (10) shows that $A[\tau^4_{\vec{\pi}}]$ is a symmetric matrix.

The ℓth matrix power of $A[\tau^4_{\vec{\pi}}]$ counts walks of length 2ℓ and type $\vec{\pi}$ by considering the form of matrix multiplication in (8). Indeed, the product $A[\tau^4_{\vec{\pi}}]^{\ell-1} A[\tau^4_{\vec{\pi}}]$ sums over the standard basis of $W^{\otimes 6}$, which can be considered as summing over the possible 6-vertex tuples used for the internal attach points of a walk of length $2(\ell-1)$ and a step. Finally, the trace of $A[\tau^4_{\vec{\pi}}]^\ell$ then counts the number of walks of length 2ℓ with both attach tuples identified and these are exactly the circuits of length 2ℓ.

6.3 Products and powers of multilinear maps

In this section, we give the formal definitions of $\phi^{2^{t-1}}$ and $A[\phi^{2^{t-1}}]$ for a t-linear map ϕ needed for the proof of Proposition 13, which appears in the next section.

Definition. *Let $V_1, \ldots, V_t$ be finite dimensional vector spaces over $\mathbb{R}$ and let $\phi, \psi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be t-linear maps. The* product *of ϕ and ψ, written $\phi * \psi$, is a $(t-1)$-linear map defined as follows. Let $u_1, \ldots, u_{t-1}$ be vectors where $u_i \in V_i$. Let $\{b_1, \ldots, b_{\dim(V_t)}\}$ be any orthonormal basis of V_t.*

$$\phi * \psi : (V_1 \otimes V_1) \times (V_2 \otimes V_2) \times \cdots \times (V_{t-1} \otimes V_{t-1}) \to \mathbb{R}$$
$$\phi * \psi(u_1 \otimes v_1, \ldots, u_{t-1} \otimes v_{t-1}) := \sum_{j=1}^{\dim(V_t)} \phi(u_1, \ldots, u_{t-1}, b_j)\psi(v_1, \ldots, v_{t-1}, b_j)$$

Extend the map $\phi * \psi$ linearly to all of the domain to produce a $(t-1)$-linear map.

Lemma 24 shows the above definition is well defined: the map is consistent when extending linearly and the map is the same for any choice of orthonormal basis $\{b_1, \ldots, b_{\dim(V_t)}\}$ of V_t.

Definition. *Let $V_1, \ldots, V_t$ be finite dimensional vector spaces over $\mathbb{R}$ and let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map and let s be an integer $0 \leq s \leq t-1$. Define*

$$\phi^{2^s} : V_1^{\otimes 2^s} \times \cdots \times V_{t-s}^{\otimes 2^s} \to \mathbb{R}$$

where $\phi^{2^0} := \phi$ and $\phi^{2^s} := \phi^{2^{s-1}} * \phi^{2^{s-1}}$.

Note that we only define this for exponents which are powers of two because the product $*$ is only defined when the domains of the maps are the same. An expression like $\phi^3 = \phi * (\phi * \phi)$ does not make sense because ϕ and $\phi * \phi$ have different domains. This defines the power $\phi^{2^{t-1}}$, which is a linear map $V_1^{\otimes 2^{t-1}} \to \mathbb{R}$. In the next section, we will prove that $\tau_{\vec{\pi}}^{2^{t-1}}$ counts the number of labeled steps of type $\vec{\pi}$, similar to (10). As we saw in Section 6.2, we need to "untangle" the attach tuples of the input to $\tau_{\vec{\pi}}^{2^{t-1}}$.

Definition. *Let $V_1, \ldots, V_t$ be finite dimensional vector spaces over $\mathbb{R}$ and let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map and define $A[\phi^{2^{t-1}}]$ to be the following square matrix/bilinear map. Let $u_1, \ldots, u_{2^{t-2}}, v_1, \ldots, v_{2^{t-2}}$ be vectors where $u_i, v_i \in V_1$.*

$$A[\phi^{2^{t-1}}] : V_1^{\otimes 2^{t-2}} \times V_1^{\otimes 2^{t-2}} \to \mathbb{R}$$

$$A[\phi^{2^{t-1}}](u_1 \otimes \cdots \otimes u_{2^{t-2}}, v_1 \otimes \ldots v_{2^{t-2}}) := \phi^{2^{t-1}}(u_1 \otimes v_1 \otimes u_2 \otimes v_2 \otimes \cdots \otimes u_{2^{t-2}} \otimes v_{2^{t-2}}).$$

Extend the map linearly to the entire domain to produce a bilinear map.

Lemma 24 from the appendix shows that the above definition is well defined, i.e. the map is consistent when extended linearly. By definition, $A[\phi^{2^{t-1}}]$ is a square real valued matrix. The next lemma, which is also proved in the appendix, shows that it is symmetric.

Lemma 16. *Let $V_1, \ldots, V_t$ be finite dimensional vector spaces over $\mathbb{R}$. If $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ is a t-linear map, then $A[\phi^{2^{t-1}}]$ is a square symmetric real valued matrix.*

6.4 Proposition 13: counting walks and circuits

This section contains the proof of Proposition 13. All of the ideas in the proof are contained in Section 6.2, which showed that $\tau_{\vec{\pi}}$, $\tau_{\vec{\pi}}^2$, and $\tau_{\vec{\pi}}^4$ counted hypergraphs, in the sense that if indicator tensors of vertices are given as input, $\tau_{\vec{\pi}}^2$ and $\tau_{\vec{\pi}}^4$ count the number of some hypergraph where the vertices from the input are fixed and all other vertices are allowed to range over all of $V(H)$. It is straightforward from the definition to see that similarly for general k and $\vec{\pi}$ and for $0 \leq s \leq t-1$, $\tau_{\vec{\pi}}^{2^s}$ will count some hypergraph. Indeed, one can iteratively expand the definition of $\tau_{\vec{\pi}}^{2^s}$ to obtain an expression of the form

$$\tau_{\vec{\pi}}^{2^s}(\cdots) = \sum \cdots \sum \underbrace{\tau_{\vec{\pi}}(\cdots) \cdots \tau_{\vec{\pi}}(\cdots)}_{2^s}. \tag{11}$$

When the inputs to $\tau_{\vec{\pi}}^{2^s}$ are indicator tensors of vertices, the above expression counts some hypergraph since the sums are over bases and each basis vector can be thought of as the indicator tensor for some vertex tuple. Thus (11) counts some hypergraph, where the specific hypergraph depends on how the inputs and summation indices are arranged on the right hand side of (11). The proof in this section comes down to showing that the hypergraph which is counted by $\tau_{\vec{\pi}}^{2^{t-1}}$ is the step $S_{\vec{\pi}}$. We do this by induction by describing exactly the hypergraph counted by $\tau_{\vec{\pi}}^{2^s}$, which is the following hypergraph.

Definition. *For $\vec{\pi} = (1, \dots, 1)$ with t parts, let $0 \leq s \leq t-1$ and define the hypergraph $D_{\vec{\pi},s}$ as follows. Let $A_1, \dots, A_{t-s}$ be disjoint sets of size 2^s where elements are labeled by binary strings of length s and let $B_{t-s+1}, \dots, B_t$ be disjoint sets of size 2^{s-1} where elements are labeled by binary strings of length $s-1$. The vertex set of $D_{\vec{\pi},s}$ is $A_1 \dot\cup \cdots \dot\cup A_{t-s} \dot\cup B_{t-s+1} \dot\cup \cdots \dot\cup B_t$. Make $a_1, \dots, a_{t-s}, b_{t-s+1}, \dots, b_t$ an edge of $D_{\vec{\pi},s}$ if $a_i \in A_i$, $b_j \in B_j$, the codes for $a_1, \dots, a_{t-s}$ are all equal, and the code for b_{t-s+j} is equal to the code formed by removing the jth bit of the code for a_1.*

For a general $\vec{\pi} = (k_1, \dots, k_t)$, start with $D_{(1,\dots,1),s}$ and expand each vertex into the appropriate size; that is, a vertex in A_i is expanded into k_i vertices and each vertex in B_j is expanded into k_j vertices. In $D_{\vec{\pi},s}$, each vertex in A_i is labeled by a pair (c, z) where c is a bit string of length s and $z \in [k_i]$. We call z the expansion index *of the vertex.*

The hypergraph $D_{\vec{\pi},0}$ is a single edge and the hypergraph $D_{\vec{\pi},t-1}$ is by definition the step $S_{\vec{\pi}}$. The following lemma precisely formulates what we mean when we say that $\tau_{\vec{\pi}}^{2^s}$ counts the hypergraph $D_{\vec{\pi},s}$.

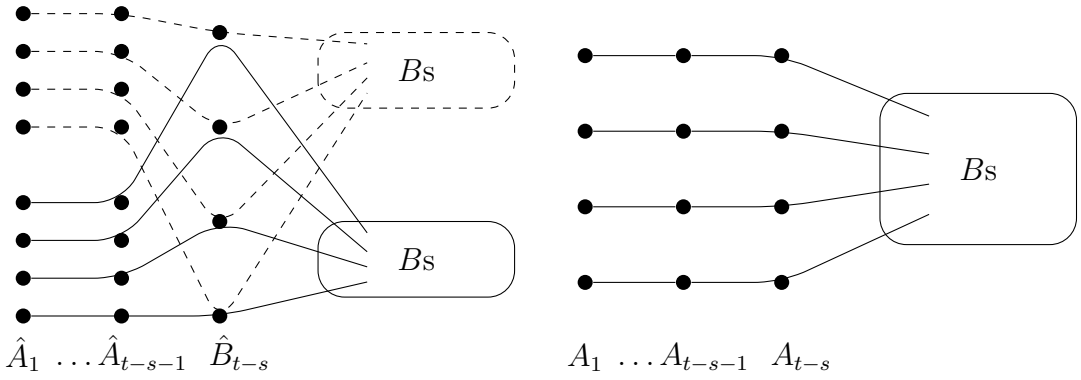
Lemma 17. *Let H be a k-uniform hypergraph, $\vec{\pi}$ a proper ordered partition of k with $\vec{\pi} = (k_1, \dots, k_t)$, and let $0 \leq s \leq t-1$. Let W be the vector space over $\mathbb{R}$ of dimension $|V(H)|$ and let τ be the adjacency map of H. Let $A_1, \dots, A_{t-s}, B_{t-s+1}, \dots, B_t$ be the vertex sets in the definition of $D_{\vec{\pi},s}$ and let Δ be any map $A_1 \cup \cdots \cup A_{t-s} \to V(H)$. Then $\tau_{\vec{\pi}}^s$ counts the number of labeled, possibly degenerate copies of $D_{\vec{\pi},s}$ extending Δ as follows.*

Let $a_{i,1}, \dots, a_{i,k_i 2^s}$ be the vertices of A_i ordered first lexicographically by bit code and then for equal codes ordered by expansion index. Let χ_i be the indicator tensor in $W^{\otimes k_i 2^s}$ for the vertex tuple $(\Delta(a_{i,1}), \dots, \Delta(a_{i,k_i 2^s}))$. Then $\tau_{\vec{\pi}}^{2^s}(\chi_1, \dots, \chi_{t-s})$ is the number of edge-preserving maps $V(D_{\vec{\pi},s}) \to V(H)$ which are consistent with Δ.

Proof. By induction on s. The base case is $s = 0$, where $D_{\vec{\pi},0}$ is a single edge, there are no B-type sets, and thus Δ is a map $V(D_{\vec{\pi},0}) \to V(H)$. The number of edge preserving maps extending Δ is either zero or one depending on if the image of Δ is an edge of H or not. But $\tau_{\vec{\pi}}(\chi_1, \dots, \chi_t)$ equals zero or one depending on if the vertices defining the indicator tensors χ_i form an edge, exactly what is required.

Assume the lemma is true for s; we will prove it for $s+1$. Denote by $\hat{A}_1, \dots, \hat{A}_{t-s-1}, \hat{B}_{t-s}, \dots, \hat{B}_t$ the sets in the definition of $D_{\vec{\pi},s+1}$ and $A_1, \dots, A_{t-s}, B_{t-s+1}, \dots, B_t$ the sets in the definition of $D_{\vec{\pi},s}$. Let $\hat{\Delta}$ be a map $\hat{A}_1 \cup \cdots \cup \hat{A}_{t-s-1} \to V(H)$ and let $\hat{\chi_1}, \dots, \hat{\chi}_{t-s-1}$ be the indicator tensors for the image of $\hat{\Delta}$ ordered as in the statement of the lemma. Since $\hat{\chi}_i$ is an indicator tensor in $W^{\otimes k_i 2^{s+1}}$, it is a simple tensor so $\hat{\chi}_i = \chi_i \otimes \chi_i'$ for $\chi_i, \chi_i' \in W^{\otimes k_i 2^s}$.

Note that χ_i is the indicator tensor for the image under $\hat{\Delta}$ of the vertices of $D_{\vec{\pi},s+1}$ whose code starts with zero and χ'_i is the indicator tensor for the image under $\hat{\Delta}$ of the vertices whose code starts with a one, since the definition of $\hat{\chi}_i$ sorted the vertices in the image lexicographically.



(a) $D_{\vec{\pi},s+1}$ (b) $D_{\vec{\pi},s}$

$$\tau_{\vec{\pi}}^{2^{s+1}}(\chi_1 \otimes \chi'_1, \ldots, \chi_{t-s-1} \otimes \chi'_{t-s-1}) = \sum_{j=1}^{d} \tau_{\vec{\pi}}^{2^s}(\chi_1, \ldots, \chi_{t-s-1}, w_j)\tau_{\vec{\pi}}^{2^s}(\chi'_1, \ldots, \chi'_{t-s-1}, w_j) \tag{12}$$

Figure 8: The induction step of Lemma 17

Consider the expansion of the definition of $\tau_{\vec{\pi}}^{2^{s+1}}(\hat{\chi}_1, \ldots, \hat{\chi}_{t-s-1})$ shown in (12) in Figure 8; the tensors $\hat{\chi}_i$ are split into χ_i and χ'_i and we sum over the standard basis $\{w_1, \ldots, w_d\}$ of $W^{\otimes k_{t-s}2^s}$, where $d = \dim(W^{\otimes k_{t-s}2^s})$. This is a generalization of Section 6.2 where (9) was compared with Figure 6 and (10) was compared with Figure 7, where Figure 6 corresponds to Figure 8 *(b)* and Figure 7 corresponds to Figure 8 *(a)*. We can consider the tensor w_j in (12) to be the indicator tensor of a tuple of $k_{t-s}2^s$ vertices.

Definition. *We now describe two embeddings of $D_{\vec{\pi},s}$ into $D_{\vec{\pi},s+1}$. In Figure 8 (a), these two embeddings are the dotted and solid lines. Let $\Gamma_0 : V(D_{\vec{\pi},s}) \to V(D_{\vec{\pi},s+1})$ be the following injection. For $1 \leq i \leq t-s-1$ and $a \in A_i$, set $\Gamma_0(a)$ equal to the vertex in $\hat{A}_i$ whose code equals the code for a with a zero prepended to the code and the same expansion index. That is, a vertex in A_i with label $(1011, 4)$ is mapped to the vertex in $\hat{A}_i$ with label $(01011, 4)$. For $a \in A_{t-s}$, set $\Gamma_0(a)$ equal to the vertex in $\hat{B}_{t-s}$ which has the same label as a. For $t-s+1 \leq j \leq t$ and $b \in B_j$, set $\Gamma_0(b)$ equal to the vertex in $\hat{B}_j$ whose code equals the code for b with a zero prepended to the code and the same expansion index. In other words, Γ_0 adds a zero to the front of the codes except for vertices in A_{t-s} whose code does not change. Define $\Gamma_1 : V(D_{\vec{\pi},s}) \to V(D_{\vec{\pi},s+1})$ similarly except prepend a one instead of a zero. In Figure 8 (a), the dotted lines represent Γ_0 and the solid lines represent Γ_1.*

Claim: Γ_0 and Γ_1 are edge preserving injections and every edge in $D_{\vec{\pi},s+1}$ is in the image of Γ_0 or Γ_1 but not both.

Proof of Claim. Let E be an edge in $D_{\vec{\pi},s}$. For $1 \leq i \leq j \leq t-s-1$ and $a_i \in A_i \cap E$ and $a_j \in A_j \cap E$, since E is an edge of $D_{\vec{\pi},s}$ the code for a_i equals the code for a_j. This implies that the codes for $\Gamma_0(a_i)$ and $\Gamma_0(a_j)$ are equal since both had a zero prepended. Now consider $b \in A_{t-s} \cap E$ which is mapped to $\hat{B}_{t-s}$. The conditions for $\Gamma_0(E)$ an edge of $D_{\vec{\pi},s+1}$ requires that the code for $\Gamma_0(b)$ equals the code formed by deleting the first bit of $\Gamma_0(a)$ where $a \in A_1 \cap E$. But the code for a equals the code for b since both are in A-type sets in $D_{\vec{\pi},s}$ and the map Γ_0 adds a zero to the front of the code for a and leaves the code for b alone. Thus the code for $\Gamma_0(b)$ equals the code formed by deleting the first bit of $\Gamma_0(a)$. Lastly, consider $b \in B_j \cap E$ for $t-s+1 \leq j \leq t$ and consider deleting the $(j+1)$-th bit of the code for $\Gamma_0(a)$. This is the same as deleting the j-th bit of a since $\Gamma_0(a)$ had a zero prepended. But deleting the jth bit of a equals the code for b, since $a, b \in E \in E(D_{\vec{\pi},s})$. Thus deleting the $(j+1)$th bit of $\Gamma_0(a)$ is the code for $\Gamma_0(b)$. We have now checked all the conditions, so $\Gamma_0(E)$ is an edge of $D_{\vec{\pi},s+1}$, i.e. Γ_0 is edge preserving. Γ_1 is edge preserving by the same argument. Finally, let E be an edge of $D_{\vec{\pi},s+1}$ and pick $a \in E \cap \hat{A}_1$. If the first bit of the code for a equals zero, then E is in the image of Γ_0 and if the first bit of the code for a equals one, then E is in the image of Γ_1. This concludes the proof of the claim. □

This claim implies that any edge-preserving map extending $\hat{\Delta}$ is formed from two edge preserving maps $V(D_{\vec{\pi},s}) \to V(H)$ each extending the appropriate restriction of $\hat{\Delta}$. Start with $\hat{\Delta}$ and extend arbitrarily to a map $\Lambda : \hat{A}_1 \cup \cdots \cup \hat{A}_{t-s-1} \cup \hat{B}_{t-s} \to V(H)$. Next define Λ_0 and Λ_1 as maps $A_1 \cup \cdots \cup A_{t-s} \to V(H)$ such that $\Lambda_0 = \Lambda \circ \Gamma_0|_{\bar{A}}$ and $\Lambda_1 = \Lambda \circ \Gamma_1|_{\bar{A}}$, where $\bar{A} = A_1 \cup \cdots \cup A_{t-s}$ so $\Gamma_0|_{\bar{A}}$ is the map Γ_0 restricted to the A-type sets in $D_{\vec{\pi},s}$. By the claim, the number of edge-preserving maps extending $\hat{\Delta}$ equals the sum over Λ of the product of the number of edge-preserving maps extending Λ_0 and extending Λ_1. This is because any edge preserving map extending Λ can be composed with Γ_0 and Γ_1 to create edge preserving maps extending Λ_0 and Λ_1, and since Γ_0 and Γ_1 are injections covering all edges of $D_{\vec{\pi},s+1}$, this can be reversed. The last step in the proof is to show that this is exactly what (12) counts.

Let $\hat{b}_1, \ldots, \hat{b}_{k_{t-s}2^s}$ be the vertices of $\hat{B}_{t-s}$ listed first in lexicographic order of codes and then by expansion index. Let w be the indicator tensor in $W^{\otimes k_{t-s}2^s}$ for the vertex tuple $(\Lambda(\hat{b}_1), \ldots, \Lambda(\hat{b}_{k_{t-s}2^s}))$. Note that as Λ ranges over all possible extensions of $\hat{\Delta}$, w ranges over the standard basis of $W^{\otimes k_{t-s}2^s}$. Now $\chi_1, \ldots, \chi_{t-s-1}, w$ are the indicator tensors representing the image of the map Λ_0, since as mentioned above, $\chi_1, \ldots, \chi_{t-s-1}$ are the indicator tensors for the image under $\hat{\Delta}$ of the vertices whose code stars with a zero. Similarly, $\chi'_1, \ldots, \chi'_{t-s-1}, w$ are the indicator tensors representing the image of the map Λ_1. Thus by induction, $\tau_{\vec{\pi}}^{2^s}(\chi_1, \ldots, \chi_{t-s-1}, w)$ is the number of edge-preserving maps extending Λ_0 and $\tau_{\vec{\pi}}^{2^s}(\chi'_1, \ldots, \chi'_{t-s-1}, w)$ is the number of edge preserving maps extending Λ_1. By the claim, this implies that the product

$$\tau_{\vec{\pi}}^{2^s}(\chi_1, \ldots, \chi_{t-s-1}, w)\tau_{\vec{\pi}}^{2^s}(\chi'_1, \ldots, \chi'_{t-s-1}, w)$$

counts the number of edge-preserving maps extending Λ. Thus (12) sums over the choices for Λ extending $\hat{\Delta}$ of the number of edge-preserving maps extending Λ. This sum is exactly the number of edge-preserving maps extending $\hat{\Delta}$, so the proof is complete. □

Corollary 18. *Let H be a k-uniform hypergraph, $\vec{\pi}$ a proper ordered partition of k with $\vec{\pi} = (k_1, \ldots, k_t)$, and let $\ell \geq 2$ be an integer. Let W be the vector space over $\mathbb{R}$ of dimension $|V(H)|$ and let τ be the adjacency map of H. Let $a_1, \ldots, a_{k_1 2^{t-2}}, a'_1, \ldots, a'_{k_1 2^{t-2}}$ be (not necessarily distinct) vertices of H and let ξ and ξ' be the indicator tensors in $W^{k_1 2^{t-2}}$ for the tuples $(a_1, \ldots, a_{k_1 2^{t-2}})$ and $(a'_1, \ldots, a'_{k_1 2^{t-2}})$ respectively. Then $A[\tau_{\vec{\pi}}^{2^{t-1}}](\xi, \xi')$ is the number of labeled, possibly degenerate steps of type $\vec{\pi}$ in H with attach tuples $(a_1, \ldots, a_{k_1 2^{t-2}})$ and $(a'_1, \ldots, a'_{k_1 2^{t-2}})$. Also, $A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell(\xi, \xi')$ is the number of labeled walks of length 2ℓ and type $\vec{\pi}$ with attach tuples $(a_1, \ldots, a_{k_1 2^{t-2}})$ and $(a'_1, \ldots, a'_{k_1 2^{t-2}})$.*

Proof. The proof is by induction on ℓ. First, consider the base case of $\ell = 1$, where the path of length two and type $\vec{\pi}$ is the step of type $\vec{\pi}$. Let A be the vertex set from the definition of the step $S_{\vec{\pi}}$. Define a mapping $\Delta : A \to V(H)$ by mapping the attach tuples of $S_{\vec{\pi}}$ to the tuples $(a_1, \ldots, a_{k_1 2^{t-2}})$ and $(a'_1, \ldots, a'_{k_1 2^{t-2}})$ in $V(H)$. By definition, the first attach tuple of $S_{\vec{\pi}}$ is the vertices ending with a zero and listed in lexicographic order and the second attach tuple of $S_{\vec{\pi}}$ is the vertices ending with a one and listed in lexicographic order. This implies that the indicator tensor χ_1 from the statement of Lemma 17 is the indicator tensor in $W^{\otimes k_1 2^{t-1}}$ for the tuple $(a_1, \ldots, a_{k_1}, a'_1, \ldots, a'_{k_1}, a_{k_1+1}, \ldots, a_{2k_1}, a'_{k_1+1}, \ldots, a'_{2k_1}, \ldots, a_{k_1 2^{t-3}+1}, \ldots, a_{k_1 2^{t-2}}, a'_{k_1 2^{t-3}+1}, \ldots, a'_{k_1 2^{t-2}})$, since each attach tuple is in lexicographic order but the last bit is zero or one so the full ordering alternates between attach tuples. By the definition of $A[\tau_{\vec{\pi}}^{2^{t-1}}]$ and the indicator tensors ξ, ξ', χ_1, $A[\tau_{\vec{\pi}}^{2^{t-1}}](\xi, \xi') = \tau_{\vec{\pi}}^{2^{t-1}}(\chi_1)$. Thus Lemma 17 applied with $s = t - 1$ shows that the number of edge-preserving maps extending Δ is $A[\tau_{\vec{\pi}}^{2^{t-1}}](\xi, \xi')$, but by the definition of Δ, this is exactly the number of labeled, possibly degenerate steps of type $\vec{\pi}$ with attach tuples $(a_1, \ldots, a_{k_1 2^{t-2}})$ and $(a'_1, \ldots, a'_{k_1 2^{t-2}})$.

Next assume that the corollary is true for ℓ; we will show that it is true for $\ell + 1$. Using the form of matrix multiplication in (8), let $\{d_1, \ldots, d_{\dim(W^{\otimes k_1 2^{t-2}})}\}$ be the standard basis of $W^{\otimes k_1 2^{t-2}}$ so

$$A[\tau_{\vec{\pi}}^{2^{t-1}}]^{\ell+1}(\xi, \xi') = \sum_{i=1}^{\dim(W^{\otimes k_1 2^{t-2}})} A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell(\xi, d_i) A[\tau_{\vec{\pi}}^{2^{t-1}}](d_i, \xi'). \tag{13}$$

Each standard basis vector d_i can be thought of as a $k_1 2^{t-2}$-tuple of vertices which corresponds to one of the two attach tuples. Thus (13) sums over the internal attach tuple for a walk of length 2ℓ and $S_{\vec{\pi}}$. □

Proof of Proposition 13. Since $A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell$ counts the number of walks of length 2ℓ, the trace counts circuits. If $\{d_1, \ldots, d_{\dim(W^{\otimes k_1 2^{t-2}})}\}$ is any orthonormal basis of $W^{\otimes k_1 2^{t-2}}$, the trace of the matrix $A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell$ is

$$\operatorname{Tr}\left[A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell\right] = \sum_{i=1}^{\dim(W^{\otimes k_1 2^{t-2}})} A[\tau_{\vec{\pi}}^{2^{t-1}}]^\ell(d_i, d_i).$$

If $\{d_1, \ldots, d_{\dim(W^{\otimes k_1 2^{t-2}})}\}$ is the standard basis, each d_i corresponds to a tuple of $k_1 2^{t-2}$ vertices, so the above expression is the number of walks of type $\vec{\pi}$ with both attach tuples equal to d_i. □

7 Applications of Hypergraph Eigenvalues

In this section, we prove Theorem 7 on Sidorenko's Conjecture [63] and prove Theorem 8, showing that there exists a polynomial-time algorithm certifying hypergraph quasirandomness.

Proof of Theorem 7. Let $k \geq 2$, let $\pi = k_1 + \cdots + k_s$ be a proper partition of k, let $\ell \geq 1$, and let G be a k-uniform hypergraph with n-vertices and e edges. Then by definition of $t_H(G)$,

$$t_{K_k}(G) = \frac{k!e}{n^k} \qquad \text{and} \qquad t_{C_{\pi,4\ell}}(G) = \frac{\#\{\text{labeled, possibly degenerate } C_{\pi,4\ell} \text{ in } G\}}{n^{|V(C_{\pi,4\ell})|}}.$$

Let $m = |E(C_{\pi,4\ell})|$, so $|V(C_{\pi,4\ell})| = mk/2$ since the cycle is two regular. We would like to prove that $t_{C_{\pi,4\ell}}(G) \geq t_{K_k}(G)^m$, which is equivalent to

$$\#\{\text{labeled, possibly degenerate } C_{\pi,4\ell} \text{ in } G\} \geq \left(\frac{k!e}{n^k}\right)^m n^{mk/2} = \left(\frac{k!e}{n^{k/2}}\right)^m.$$

By Proposition 13, the number of labeled, possibly degenerate $C_{\pi,4\ell}$ in G is the trace of $A[\tau_{\vec{\pi}}^{2^{s-1}}]^{2\ell}$. Let $A = A[\tau_{\vec{\pi}}^{2^{s-1}}]$, let $\hat{1}$ be the all-ones vector scaled to unit length, and let $\mu_1, \ldots, \mu_d$ be the eigenvalues of A arranged so that $|\mu_1| \geq \cdots \geq |\mu_d|$. Then since $\mu_i^{2\ell} \geq 0$ and using Lemma 15 and that $m = |E(C_{\pi,4\ell})| = 2\ell 2^{s-1}$, we have that

$$\left(\frac{k!e}{n^{k/2}}\right)^m = \left(\tau_{\vec{\pi}}(\hat{1}, \ldots, \hat{1})\right)^m \leq \|\tau_{\vec{\pi}}\|^m \leq \mu_1^{2\ell} \leq \operatorname{Tr}\left[A^{2\ell}\right],$$

completing the proof. □

We now turn our attention to Theorem 8, which is proved by the following algorithm.

Algorithm 1. *On input $k \geq 4$, $\epsilon, \eta > 0$, and a k-uniform hypergraph H with m edges and n vertices,*

- *Let $\pi = \lfloor k/2 \rfloor + \lceil k/2 \rceil$.*
- *Define ℓ as follows:*

$$\ell = \ell_\epsilon = \begin{cases} \left\lceil \frac{k}{4\epsilon} \right\rceil & \text{if } k \text{ is even,} \\ \left\lceil \frac{k+1}{4\sqrt{k}-2} \right\rceil & \text{otherwise.} \end{cases} \tag{14}$$

- *Let N be the number of labeled circuits of type π and length 4ℓ in H.*

- *If* $N < (1+\eta)(k!)^{4\ell} m^{4\ell} n^{-2\ell k}$ *then output* ***Quasirandom****, otherwise output* ***Unknown****.*

Just counting labeled circuits directly, Algorithm 1 runs in time $O(n^{2k\ell})$ since $C_{\pi,4\ell}$ has $2k\ell$ vertices. For k even this is $n^{O(k^2)}$ and for k odd is $n^{O(k^3/2)}$. Using Proposition 13, an alternate way of counting labeled circuits is to compute the trace of $A[\tau_{\vec{\pi}}^2]^{2\ell}$. The matrix $A = A[\tau_{\vec{\pi}}^2]$ can be computed in time $O(n^{2k} \operatorname{polylog} n)$ since there are $n^{2\lfloor k/2 \rfloor}$ basis vectors on which to evaluate $\tau_{\vec{\pi}} * \tau_{\vec{\pi}}$ and each evaluation takes $O(n^{\lceil k/2 \rceil})$ multiplications. Note that the maximum size of entries of A is $n^{2k\ell}$. Since A is an $n^{2\lfloor k/2 \rfloor} \times n^{2\lfloor k/2 \rfloor}$-matrix, using repeated squaring we can compute $A^{2\ell}$ in time $O(n^{k\omega} \log \ell \operatorname{polylog} n)$, where ω is the exponent for matrix multiplication. Thus Algorithm 1 runs in time $O(n^{k\omega} \operatorname{polylog} n)$.

Theorem 19. *Let* $k \geq 4$ *and* $\epsilon, \delta > 0$ *be given. There exists a constant* $\eta > 0$ *depending only on* δ *and* k *such that the following holds. If on input* k*,* ϵ*,* η*, and a* k*-uniform,* n*-vertex,* m*-edge hypergraph* H *Algorithm 1 outputs* ***Quasirandom****, then for all* $S_1 \subseteq \binom{V(H)}{\lceil k/2 \rceil}$ *and all* $S_2 \subseteq \binom{V(H)}{\lfloor k/2 \rfloor}$*,*

$$\left| e(S_1, S_2) - \frac{k!m}{n^k} |S_1||S_2| \right| \leq \frac{\delta m}{n^{k/2}} \sqrt{|S_1||S_2|}. \tag{15}$$

In addition, if

$$p \geq \begin{cases} n^{-k/2+\epsilon} & \text{if } k \text{ is even,} \\ n^{-k/2+\sqrt{k}} & \text{otherwise,} \end{cases} \tag{16}$$

then with probability going to one as n *goes to infinity, for a hypergraph* H *drawn from the distribution* $G^{(k)}(n,p)$*, Algorithm 1 on input* k*,* ϵ*,* η*, and* H *outputs* ***Quasirandom****.*

Proof of the correctness of Algorithm 1. We prove that if Algorithm 1 outputs `Quasirandom` then the expansion property (15) holds. Let k, ϵ, δ, π, and ℓ be defined as in the theorem and let H be any k-uniform hypergraph with m edges and n vertices. We need to prove that there exists an $\eta > 0$ so that if the number of labeled circuits of type π and length 4ℓ is at most $(1+\eta)(k!)^{4\ell} m^{4\ell} n^{-2\ell k}$, then (15) holds. This proof is very similar to the proof that $\texttt{Cycle}_{4\ell}[\pi] \Rightarrow \texttt{Eig}[\pi]$ on page 24 and we use the same notation. Let $\hat{1}$ denote the all-ones vector scaled to unit length, let $A = A[\tau_{\vec{\pi}}^2]$, and let $\mu_1, \ldots, \mu_d$ be the eigenvalues of A arranged so that $|\mu_1| \geq \cdots \geq |\mu_d|$. Since $\mu_i^{2\ell} \geq 0$, Proposition 13 implies that

$$\mu_1^{2\ell} \leq \mu_1^{2\ell} + \mu_2^{2\ell} \leq \operatorname{Tr}\left[A^{2\ell}\right] = \#\{\text{possibly degenerate } C_{\pi,4\ell} \text{ in } H\} \leq (1+\eta) \frac{(k!)^{4\ell} m^{4\ell}}{n^{2\ell k}}. \tag{17}$$

Using Lemma 15 and (17), we have that

$$\frac{k!m}{n^{k/2}} = \tau_{\vec{\pi}}(\hat{1}, \ldots, \hat{1}) \leq \|\tau_{\vec{\pi}}\| \leq \sqrt{\mu_1} \leq (1+\eta)^{1/4\ell} \frac{k!m}{n^{k/2}}.$$

Combining this with (17), we obtain

$$\frac{k!^2 m^2}{n^k} \le \mu_1 \le (1+\eta)^{1/2\ell} \frac{k!^2 m^2}{n^k} \qquad \text{and} \qquad \mu_2 \le \eta^{1/2\ell} \frac{k!^2 m^2}{n^k}.$$

Since $\tau_{\vec{\pi}}(\hat{1},\ldots,\hat{1})^2 = \frac{k!^2 m^2}{n^k}$, Proposition 14 implies that given $\delta > 0$, it is possible to choose $\eta > 0$ such that $\lambda_{2,\pi}(H) = \|\tau_{\vec{\pi}} - qJ_{\vec{\pi}}\| \le \delta \frac{m}{n^{k/2}}$. Finally, the Hypergraph Expander Mixing Lemma (Theorem 10) shows that for all $S_1 \subseteq \binom{V(H)}{\lceil k/2 \rceil}$ and all $S_2 \subseteq \binom{V(H)}{\lfloor k/2 \rfloor}$

$$\left| e(S_1,S_2) - \frac{k!m}{n^k}|S_1||S_2| \right| \le \lambda_{2,\pi}(H)\sqrt{|S_1||S_2|} \le \frac{\delta m}{n^{k/2}}\sqrt{|S_1||S_2|}.$$

□

All that remains to prove Theorem 19 is to prove that for almost all k-uniform hypergraphs with probability satisfying (16), Algorithm 1 outputs `Quasirandom`. This is proved in two steps; let G be a hypergraph drawn from the distribution $G^{(k)}(n,p)$. We first prove that the expected number of labeled circuits in G is $(1+o(1))p^{4\ell}n^{2k\ell}$ (note that $C_{\pi,4\ell}$ has 4ℓ edges and $2k\ell$ vertices, but the proof is complicated by the presence of degenerate cycles). Second, we use the second moment method to prove that with high probability the number of labeled circuits in G is close to the expectation. This in turn will imply that w.h.p. Algorithm 1 outputs `Quasirandom`, since w.h.p. $|E(G)| = (1+o(1))p\binom{n}{k}$ so that w.h.p. $(1+o(1))p^{4\ell}n^{2k\ell} = (1+o(1))(k!)^{4\ell}|E(G)|^{4\ell}n^{-2k\ell}$, which is the bound checked by Algorithm 1. A *degenerate cycle of type π and length 4ℓ* is a k-uniform hypergraph H for which there exists an edge-preserving surjection $\phi : V(C_{\pi,4\ell}) \to V(H)$ where ϕ is not an injection.

Lemma 20. *Let $k \ge 2$, $\ell \ge 1$, and $\pi = \lfloor k/2 \rfloor + \lceil k/2 \rceil$. Let H be a degenerate cycle of type π and length 4ℓ with v vertices and e edges. Then*

$$e \ge \begin{cases} \frac{v}{\lceil k/2 \rceil} & \text{if } v > 2\ell \left\lceil \frac{k}{2} \right\rceil + \left\lfloor \frac{k}{2} \right\rfloor, \\ \frac{v}{\lceil k/2 \rceil} - 1 & \text{otherwise.} \end{cases}$$

Proof. Let $\phi : V(C_{\pi,4\ell}) \to V(H)$ be an edge-preserving surjection. Let $E_1,\ldots,E_{4\ell}$ be the edges of $C_{\pi,4\ell}$, so that $\phi(E_1),\ldots,\phi(E_{4\ell})$ are the edges of H (possibly with repetitions). Since ϕ is edge-preserving, $|\phi(E_1) \cap \phi(E_2)| \ge \lfloor k/2 \rfloor$ so if $\phi(E_2)$ is a distinct edge from $\phi(E_1)$ then $\phi(E_2)$ has at most $\lceil k/2 \rceil$ vertices outside $\phi(E_1)$. Similarly, if $\phi(E_3)$ is an edge of H distinct from $\phi(E_1)$ and $\phi(E_2)$ then it can have at most $\lceil k/2 \rceil$ vertices of H outside $\phi(E_1) \cup \phi(E_2)$ since $\phi(E_3)$ must intersect $\phi(E_2)$ in at least $\lfloor k/2 \rfloor$ vertices. In general, if $\phi(E_i)$ is an edge distinct from $\phi(E_1),\ldots,\phi(E_{i-1})$ then it can have at most $\lceil k/2 \rceil$ vertices of H outside $\cup_{j<i}\phi(E_j)$. Since H has no isolated vertices (ϕ is an edge-preserving surjection),

$$v = |V(H)| = \left| \cup \phi(E_i) \right| \le |\phi(E_1)| + (e-1)\left\lceil \frac{k}{2} \right\rceil = k + (e-1)\left\lceil \frac{k}{2} \right\rceil \tag{18}$$

$$\Leftrightarrow \quad v - \lfloor k/2 \rfloor \le e \lceil k/2 \rceil$$

$$\Leftrightarrow \quad \frac{v}{\lceil k/2 \rceil} - 1 \le \frac{v - \lfloor k/2 \rfloor}{\lceil k/2 \rceil} \le e.$$

Now assume that $v > 2\ell \left\lceil \frac{k}{2} \right\rceil + \left\lfloor \frac{k}{2} \right\rfloor$. We show that the lower bound on e from the previous paragraph can be improved by one. First, assume that for all i, there exists a $j \neq i$ such that $\phi(E_i) = \phi(E_j)$. Since each edge of H has at least two edges of $C_{\pi,4\ell}$ mapped to it, $e = |E(H)| \leq \frac{1}{2}|E(C_{\pi,4\ell})| = 2\ell$. Define disjoint vertex sets $A_1, \ldots, A_{4\ell} \subseteq V(C_{\pi,4\ell})$ such that $E_i = A_i \cup A_{i+1}$ for $i < 4\ell$ and $E_{4\ell} = A_{4\ell} \cup A_1$. Now define an auxiliary graph G as follows. First, set $V(G) = \{\phi(A_i) : 1 \leq i \leq 4\ell\}$. (Note this is somewhat subtle: for each set $\phi(A_i)$ we create a vertex of G, except if for some $i \neq j$ we have $\phi(A_i) = \phi(A_j)$, in which case the same vertex of G is used. On the other hand, if $\phi(A_i) \cap \phi(A_j) \neq \emptyset$ but $\phi(A_i) \neq \phi(A_j)$, then separate vertices of G are created.) For each hyperedge $\phi(E_i)$ of H with $i < 4\ell$, add a graph edge between $\phi(A_i)$ and $\phi(A_{i+1})$; in addition, add an edge between $\phi(A_{4\ell})$ and $\phi(A_1)$. The number of edges of G is the same as the number of hyperedges of H so $|E(G)| = e \leq 2\ell$. Also, G is connected since H is a degenerate cycle. Since G is connected with at most 2ℓ edges, G has at most $2\ell + 1$ vertices. Each vertex of G translates into at most $\lceil k/2 \rceil$ vertices of H, but if k is odd there must be at least one vertex of G representing a set of size $\lfloor k/2 \rfloor$. Thus $|V(H)| \leq 2\ell \lceil k/2 \rceil + \lfloor k/2 \rfloor$, which contradicts our assumption that $v > 2\ell \lceil k/2 \rceil + \lfloor k/2 \rfloor$.

Thus we can assume that there exists some E_i such that for all $j \neq i$, $\phi(E_i) \neq \phi(E_j)$. By symmetry, relabel the edges of the cycle so that $i = 1$, i.e. for all $j \neq 1$, $\phi(E_1) \neq \phi(E_j)$. One of the two (or both) of the following occur:

- $\phi(A_1) \subseteq \bigcup_{i>1} \phi(A_i)$,
- $\phi(E_{4\ell})$ is distinct as an edge of H from $\phi(E_1), \ldots, \phi(E_{4\ell-1})$.

Indeed, assume that $\phi(A_1)$ is not contained in $\cup_{i>1}\phi(A_i)$. Since $\phi(E_{4\ell})$ includes $\phi(A_1)$ and $\phi(E_2) \cup \cdots \cup \phi(E_{4\ell-1}) = \cup_{i>1}\phi(A_i)$, $\phi(E_{4\ell})$ is distinct from $\phi(E_2), \ldots, \phi(E_{4\ell-1})$. But by assumption, $\phi(E_1) \neq \phi(E_{4\ell})$, so that $\phi(E_{4\ell})$ is distinct as a hyperedge of H.

Assume that $\phi(A_1) \subseteq \cup_{i>1}\phi(A_i)$, and consider the argument from the first paragraph: consider edges of the cycle one by one starting from E_1. Each $\phi(E_j)$ which is distinct can add at most $\lceil k/2 \rceil$ new vertices to the union $\cup_{1<i\leq j}\phi(A_i)$, since $\phi(E_j)$ must share at least $\lfloor k/2 \rfloor$ with $\phi(E_{j-1})$. Since $\phi(A_1)$ is a subset of $\cup_{i>1}\phi(A_i)$, the inequality in (18) can be improved to

$$v = |\cup_{i>1}\phi(A_i)| \leq e \left\lceil \frac{k}{2} \right\rceil .$$

Now assume that $\phi(E_{4\ell})$ is distinct as an edge of H from $\phi(E_1), \ldots, \phi(E_{4\ell-1})$. By definition, $\phi(E_{4\ell})$ shares $\phi(A_{4\ell})$ with $\phi(E_{4\ell-1})$ and shares $\phi(A_1)$ with $\phi(E_1)$. Thus $\phi(E_{4\ell})$ is a hyperedge of H distinct from $\phi(E_1), \ldots, \phi(E_{4\ell-1})$ which does not use any new vertices. We can therefore improve the bound on the number of edges by one. □

Lemma 21. *Let k be an even integer at least 4, $\epsilon > 0$, $\pi = k/2 + k/2$, and ℓ as defined in (14). For $p \geq n^{-k/2+\epsilon}$, the expected number of labeled, degenerate circuits of type π and length 4ℓ in $G^{(k)}(n,p)$ is $o(p^{4\ell}n^{2k\ell})$.*

Proof. We need to prove that for every degenerate cycle H of type π and length 4ℓ we have $p^{|E(H)|}n^{|V(H)|} = o(p^{4\ell}n^{2k\ell})$. Indeed, the expected number of labeled copies of H in $G^{(k)}(n,p)$

is $p^{|E(H)|}n^{|V(H)|}$ and there are constantly many degenerate cycles since ℓ and k are constants independent of n and p. Thus if $p^{|E(H)|}n^{|V(H)|} = o(p^{4\ell}n^{2k\ell})$ for all H, the expected number of degenerate cycles is $o(p^{4\ell}n^{2k\ell})$.

Let H be a degenerate cycle of type π and length 4ℓ with v vertices and e edges. We need to prove that $p^e n^v = o(p^{4\ell}n^{2k\ell})$. Substituting in p, we need to prove that

$$n^{v-\frac{ek}{2}+e\epsilon} = o(n^{2k\ell-\frac{4\ell k}{2}+4\ell\epsilon}),$$

i.e. we need to prove that

$$v - e\left(\frac{k}{2}-\epsilon\right) < 4\ell\epsilon. \tag{19}$$

Case 1: $v > \ell k + \frac{k}{2}$. By Lemma 20, $e \geq \frac{2v}{k}$. If $e = 4\ell$, then H has the same number of edges as $C_{\pi,4\ell}$ but fewer vertices so trivially $p^e n^v = o(p^{4\ell}n^{2k\ell})$. Thus $e < 4\ell$ which implies that $v \leq 2k\ell - \frac{k}{2}$, since the map $\phi : V(C_{\pi,4\ell}) \to V(H)$ must map at least two sets of size $k/2$ in $V(C_{\pi,4\ell})$ to the same vertex set in H. Inserting these bounds into (19) and simplifying, we obtain

$$v - e\left(\frac{k}{2}-\epsilon\right) \leq v - \frac{2v}{k}\left(\frac{k}{2}-\epsilon\right) = \frac{2v\epsilon}{k} \leq \frac{2\epsilon}{k}\left(2k\ell - \frac{k}{2}\right) = 4\ell\epsilon - \epsilon < 4\ell\epsilon.$$

Case 2: $v \leq \ell k + \frac{k}{2}$. By Lemma 20, we have that $e \geq \frac{2v}{k} - 1$. Inserting these bounds into (19), we obtain

$$\begin{aligned} v - e\left(\frac{k}{2}-\epsilon\right) \leq v - \left(\frac{2v}{k}-1\right)\left(\frac{k}{2}-\epsilon\right) = \frac{2v\epsilon}{k} + \frac{k}{2} - \epsilon \leq \frac{2\epsilon}{k}\left(\ell k + \frac{k}{2}\right) + \frac{k}{2} - \epsilon \\ = 2\ell\epsilon + \frac{k}{2}. \end{aligned}$$

Thus to prove (19), we require

$$2\ell\epsilon + \frac{k}{2} < 4\ell\epsilon \quad \Leftrightarrow \quad \frac{k}{4\epsilon} < \ell,$$

which is true by the definition of ℓ. □

Lemma 22. *Let k be an odd integer at least 5, $\pi = \lfloor k/2 \rfloor + \lceil k/2 \rceil$, and ℓ as defined in (14). For $p \geq n^{-k/2+\sqrt{k}}$, the expected number of labeled, degenerate cycles of type π and length 4ℓ in $G^{(k)}(n,p)$ is $o(p^{4\ell}n^{2k\ell})$.*

Proof. Similar to the previous proof, we need to prove that $p^{|E(H)|}n^{|V(H)|} = o(p^{4\ell}n^{2k\ell})$ for all degenerate cycles H. Let H be a degenerate cycle of type π and length 4ℓ with v vertices and e edges. Substituting in the value of p, we need to prove that

$$v - e\left(\frac{k}{2}-\sqrt{k}\right) < 4\ell\sqrt{k}. \tag{20}$$

Case 1: $v > 2\ell\lceil k/2\rceil + \lfloor k/2\rfloor$. In this case, Lemma 20 implies that $e \geq v/\lceil k/2\rceil$. If $e = 4\ell$, then H has the same number of edges as $C_{\pi,4\ell}$ but fewer vertices so trivially $p^e n^v = o(p^{4\ell} n^{2k\ell})$. Thus $e < 4\ell$ which implies that $v \leq 2k\ell - \lfloor k/2\rfloor$, since the map $\phi : V(C_{\pi,4\ell}) \to V(H)$ must map at least two sets of size $\lfloor k/2\rfloor$ in $V(C_{\pi,4\ell})$ to the same vertex set in H. Starting from (20), plugging in these two bounds on v and e, and using that $k \geq 5$ we obtain

$$\begin{aligned} v - e\left(\frac{k}{2} - \sqrt{k}\right) \leq v - \frac{v}{\lceil k/2\rceil}\left(\frac{k}{2} - \sqrt{k}\right) &= v\left(1 - \frac{2}{k+1}\left(\frac{k}{2} - \sqrt{k}\right)\right) \\ &\leq \left(2\ell k - \frac{k-1}{2}\right)\left(1 - \frac{k}{k+1} + \frac{2\sqrt{k}}{k+1}\right). \quad (21) \end{aligned}$$

Thus to prove (20), we require that (21) is less than $4\ell\sqrt{k}$. Solving for ℓ, we obtain

$$\left(2\ell k - \frac{k-1}{2}\right)\left(1 - \frac{k}{k+1} + \frac{2\sqrt{k}}{k+1}\right) < 4\ell\sqrt{k} \quad \Leftrightarrow \quad \ell < \frac{2k^{3/2} + k - 2\sqrt{k} - 1}{4k - 8\sqrt{k}}. \quad (22)$$

Therefore, to complete the proof we need to check that the definition of ℓ from from (14) makes this inequality true for all $k \geq 5$.

Since $8k\sqrt{k} < 4k^2$ and $8\sqrt{k} < 26k$ for $k \geq 5$,

$$8k\sqrt{k} + 8\sqrt{k} < 4k^2 + 26k + 2.$$

This implies that

$$4k^2 + 8k\sqrt{k} - 36k + 8\sqrt{k} < 8k^2 - 10k + 2$$

which implies that

$$\begin{aligned} (k + 4\sqrt{k} - 1)(4k - 8\sqrt{k}) &< (2k^{3/2} + k - 2\sqrt{k} - 1)(4\sqrt{k} - 2) \\ \frac{k+1}{4\sqrt{k} - 2} + 1 &< \frac{2k^{3/2} + k - 2\sqrt{k} - 1}{4k - 8\sqrt{k}}. \end{aligned}$$

Since $\ell = \left\lceil \frac{k+1}{4\sqrt{k}-2} \right\rceil$, (22) holds.

Case 2: $v \leq 2\ell\lceil k/2\rceil + \lfloor k/2\rfloor$. In this case, Lemma 20 implies that $e \geq v/\lceil k/2\rceil - 1$. Plugging these two bounds into (20), we obtain

$$\begin{aligned} v - e\left(\frac{k}{2} - \sqrt{k}\right) &\leq v - \left(\frac{v}{\lceil k/2\rceil} - 1\right)\left(\frac{k}{2} - \sqrt{k}\right) \\ &\leq \left(2\ell\left\lceil\frac{k}{2}\right\rceil + \left\lfloor\frac{k}{2}\right\rfloor\right)\left(1 - \frac{1}{\lceil k/2\rceil}\left(\frac{k}{2} - \sqrt{k}\right)\right) + \frac{k}{2} - \sqrt{k}. \quad (23) \end{aligned}$$

Since $1 - 1/\lceil k/2 \rceil (k/2 - \sqrt{k}) = \frac{2\sqrt{k}+1}{k+1}$, (23) continues as

$$\left(\ell(k+1) + \frac{k-1}{2}\right)\left(\frac{2\sqrt{k}+1}{k+1}\right) + \frac{k}{2} - \sqrt{k} < 2\ell\sqrt{k} + \ell + \sqrt{k} + \frac{1}{2} + \frac{k}{2} - \sqrt{k}$$
$$= 2\ell\sqrt{k} + \ell + \frac{k}{2} + \frac{1}{2}.$$

Therefore, to prove (20) we must prove that

$$2\ell\sqrt{k} + \ell + \frac{k}{2} + \frac{1}{2} \leq 4\ell\sqrt{k} \quad \Leftrightarrow \quad \ell \geq \frac{k+1}{4\sqrt{k}-2},$$

which is true by the definition of ℓ. □

Proof of Theorem 19. We need to prove that for all $\eta > 0$ and for p defined as in (16), with probability going to one as n goes to infinity, a graph G drawn from the distribution $G^{(k)}(n,p)$ has the following property: the number of labeled circuits of type π and length 4ℓ in G is at most $(1+\eta)(k!)^{4\ell}|E(G)|^{4\ell}n^{-2k\ell}$. With high probability, $|E(G)| = (1 \pm \frac{\eta}{16\ell})\frac{pn^k}{k!}$. Therefore, we need to prove that with high probability the number of labeled circuits in G is at most $(1+\frac{\eta}{2})p^{4\ell}n^{2k\ell}$.

We prove this in two stages. First, we prove that with high probability the number of labeled, degenerate cycles is upper bounded by $\frac{\eta}{4}p^{4\ell}n^{2k\ell}$. This is just the first moment method as follows. Let Y be the number of labeled, degenerate cycles in G. By Markov's Inequality,

$$\mathbb{P}\left[Y > \frac{\eta p^{4\ell} n^{2k\ell}}{4}\right] < \frac{4\,\mathbb{E}[Y]}{\eta p^{4\ell} n^{2k\ell}}.$$

Lemmas 21 and 22 prove that the right hand side goes to zero as n goes to infinity, so with high probability the number of labeled, degenerate cycles in G is at most $\frac{\eta}{4}p^{4\ell}n^{2k\ell}$.

We now use the second moment method to prove that with high probability, the number of labeled (non-degenerate) cycles is concentrated around its expectation. First by the definition of p, the expected number of cycles goes to infinity. For each list W of $2k\ell$ distinct vertices of G, define an event A_W as "W forms a labeled $C_{\pi,4\ell}$ in G". Using the techniques in [7, Section 4.3], the second moment method then comes down to showing that for a list W,

$$\Delta^* = \sum_{W' \sim W} \mathbb{P}[A_{W'}|A_W] = o(p^{4\ell}n^{2k\ell}),$$

where $W' \sim W$ means that the events A_W and $A_{W'}$ are dependent.

Assume that $\hat{W}$ and $\hat{W}'$ are cycles in $K_n^{(k)}$ which share at least one edge and let W and W' be the lists of vertices of $\hat{W}$ and $\hat{W}'$ respectively. Then we have that A_W and $A_{W'}$ are

dependent. Let v be the number of vertices of $\hat{W}'$ which do not appear in $V(\hat{W})$. Since the cycle is two regular, there must be at least $\frac{2v}{k}$ edges of $\hat{W}'$ which are not in $E(\hat{W})$. Thus

$$\mathbb{P}[A_{W'}|A_W] \leq p^{2v/k}.$$

There are at most n^v choices for v vertices outside $V(\hat{W})$, so

$$\Delta^* \leq \sum_{v=1}^{2k\ell-k} p^{2v/k} n^v. \tag{24}$$

The algebra in Case 1 of Lemmas 21 and 22 computed that if e and v are integers such that $e \geq v/\lceil k/2 \rceil$ and $v \leq 2k\ell - \lfloor k/2 \rfloor$, then $p^e n^v = o(p^{4\ell} n^{2k\ell})$. Since W and W' must share at least k vertices, $v \leq 2k\ell - k \leq 2k\ell - \lfloor k/2 \rfloor$, so if we let $e = 2v/k \geq v/\lceil k/2 \rceil$, then the computations in Case 1 of Lemmas 21 and 22 imply that $p^{2v/k} n^v = o(p^{4\ell} n^{2k\ell})$ for all $1 \leq v \leq 2k\ell - k$. Since there are constantly many terms in the sum (24), $\Delta^* = o(p^{4\ell} n^{2k\ell})$ which implies that with probability going to one as n goes to infinity, the number of cycles is at most $(1 + \frac{\eta}{4}) p^{4\ell} n^{2k\ell}$. Combined with the fact that with high probability the number of degenerate cycles is at most $\frac{\eta}{4} p^{4\ell} n^{2k\ell}$, the number of labeled circuits in G is at most $(1 + \frac{\eta}{2}) p^{4\ell} n^{2k\ell}$ with high probability, completing the proof. □

Acknowledgements. The authors would like to thank Vojtech Rödl and Mathias Schacht for helpful discussion and feedback.

References

[1] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6:83–96, 1986.

[2] N. Alon, A. Coja-Oghlan, H. Hàn, M. Kang, V. Rödl, and M. Schacht. Quasi-randomness and algorithmic regularity for graphs with general degree distributions. *SIAM J. Comput.*, 39:2336–2362, 2010.

[3] N. Alon, R. A. Duke, H. Lefmann, V. Rödl, and R. Yuster. The algorithmic aspects of the regularity lemma. *J. Algorithms*, 16:80–109, 1994.

[4] N. Alon and V. D. Milman. λ_1, isoperimetric inequalities for graphs, and superconcentrators. *J. Combin. Theory Ser. B*, 38:73–88, 1985.

[5] N. Alon and P. Pudlák. Constructive lower bounds for off-diagonal Ramsey numbers. *Israel J. Math.*, 122:243–251, 2001.

[6] N. Alon and V. Rödl. Sharp bounds for some multicolor Ramsey numbers. *Combinatorica*, 25:125–141, 2005.

[7] N. Alon and J. H. Spencer. *The probabilistic method*. Wiley-Interscience Series in Discrete Mathematics and Optimization. John Wiley & Sons Inc., Hoboken, NJ, third edition, 2008.

[8] T. Austin and T. Tao. Testability and repair of hereditary hypergraph properties. *Random Structures Algorithms*, 36:373–463, 2010.

[9] Y. Bilu and S. Hoory. On codes from hypergraphs. *European J. Combin.*, 25:339–354, 2004.

[10] F. Chung. Quasi-random hypergraphs revisited. *Random Structures Algorithms*, 40:39–48, 2012.

[11] F. Chung and R. Graham. Sparse quasi-random graphs. *Combinatorica*, 22:217–244, 2002.

[12] F. Chung and R. Graham. Quasi-random graphs with given degree sequences. *Random Structures Algorithms*, 32:1–19, 2008.

[13] F. R. K. Chung. Quasi-random classes of hypergraphs. *Random Structures Algorithms*, 1:363–382, 1990.

[14] F. R. K. Chung. The Laplacian of a hypergraph. In *Expanding graphs (Princeton, NJ, 1992)*, volume 10 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pages 21–36. Amer. Math. Soc., Providence, RI, 1993.

[15] F. R. K. Chung and R. L. Graham. Quasi-random hypergraphs. *Random Structures Algorithms*, 1:105–124, 1990.

[16] F. R. K. Chung and R. L. Graham. Quasi-random set systems. *J. Amer. Math. Soc.*, 4:151–196, 1991.

[17] F. R. K. Chung and R. L. Graham. Cohomological aspects of hypergraphs. *Trans. Amer. Math. Soc.*, 334:365–388, 1992.

[18] F. R. K. Chung, R. L. Graham, and R. M. Wilson. Quasi-random graphs. *Combinatorica*, 9:345–362, 1989.

[19] A. Coja-Oghlan, C. Cooper, and A. Frieze. An efficient sparse regularity concept. *SIAM J. Discrete Math.*, 23:2000–2034, 2009/10.

[20] A. Coja-Oghlan, A. Goerdt, A. Lanka, and F. Schädlich. Techniques from combinatorial approximation algorithms yield efficient algorithms for random $2k$-SAT. *Theoret. Comput. Sci.*, 329:1–45, 2004.

[21] D. Conlon, J. Fox, and B. Sudakov. An approximate version of Sidorenko's conjecture. *Geom. Funct. Anal.*, 20:1354–1366, 2010.

[22] D. Conlon, H. Hàn, Y. Person, and M. Schacht. Weak quasi-randomness for uniform hypergraphs. *Random Structures Algorithms*, 40:1–38, 2012.

[23] J. Cooper and A. Dutle. Spectra of uniform hypergraphs.

[24] D. Dellamonica, P. Haxell, T. Łuczak, D. Mubayi, B. Nagle, Y. Person, V. Rodl, and M. Schacht. Tree-minimal graphs are almost regular. accepted.

[25] D. Dellamonica, Jr. and V. Rödl. Hereditary quasirandom properties of hypergraphs. *Combinatorica*, 31:165–182, 2011.

[26] I. Dinur. The PCP theorem by gap amplification. *J. ACM*, 54:Art. 12, 44, 2007.

[27] P. Erdős and A. Hajnal. On Ramsey like theorems. Problems and results. In *Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972)*, pages 123–140. Inst. Math. Appl., Southend, 1972.

[28] U. Feige and E. Ofek. Easily refutable subformulas of large random 3CNF formulas. In *Automata, languages and programming*, volume 3142 of *Lecture Notes in Comput. Sci.*, pages 519–530. Springer, Berlin, 2004.

[29] K. Feng and W.-C. W. Li. Spectra of hypergraphs and applications. *J. Number Theory*, 60:1–22, 1996.

[30] J. Friedman. Some graphs with small second eigenvalue. *Combinatorica*, 15:31–42, 1995.

[31] J. Friedman, A. Goerdt, and M. Krivelevich. Recognizing more unsatisfiable random k-SAT instances efficiently. *SIAM J. Comput.*, 35:408–430, 2005.

[32] J. Friedman and A. Wigderson. On the second eigenvalue of hypergraphs. *Combinatorica*, 15:43–65, 1995.

[33] W. T. Gowers. Quasirandomness, counting and regularity for 3-uniform hypergraphs. *Combin. Probab. Comput.*, 15:143–184, 2006.

[34] W. T. Gowers. Hypergraph regularity and the multidimensional Szemerédi theorem. *Ann. of Math. (2)*, 166:897–946, 2007.

[35] W. T. Gowers. Quasirandom groups. *Combin. Probab. Comput.*, 17:363–387, 2008.

[36] H. Hàn, Y. Person, and M. Schacht. Note on strong refutation algorithms for random k-SAT formulas. In *LAGOS'09—V Latin-American Algorithms, Graphs and Optimization Symposium*, volume 35 of *Electron. Notes Discrete Math.*, pages 157–162. Elsevier Sci. B. V., Amsterdam, 2009.

[37] H. Hatami. Graph norms and Sidorenko's conjecture. *Israel J. Math.*, 175:125–150, 2010.

[38] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43:439–561 (electronic), 2006.

[39] S. Janson. Quasi-random graphs and graph limits. *European J. Combin.*, 32:1054–1083, 2011.

[40] P. Keevash. A hypergraph regularity method for generalized Turán problems. *Random Structures Algorithms*, 34:123–164, 2009.

[41] Y. Kohayakawa, B. Nagle, V. Rödl, and M. Schacht. Weak hypergraph regularity and linear hypergraphs. *J. Combin. Theory Ser. B*, 100:151–160, 2010.

[42] Y. Kohayakawa and V. Rödl. Regular pairs in sparse random graphs. I. *Random Structures Algorithms*, 22:359–434, 2003.

[43] Y. Kohayakawa, V. Rödl, and P. Sissokho. Embedding graphs with bounded degree in sparse pseudorandom graphs. *Israel J. Math.*, 139:93–137, 2004.

[44] Y. Kohayakawa, V. Rödl, and J. Skokan. Hypergraphs, quasi-randomness, and conditions for regularity. *J. Combin. Theory Ser. A*, 97:307–352, 2002.

[45] M. Krivelevich and B. Sudakov. Pseudo-random graphs. In *More sets, graphs and numbers*, volume 15 of *Bolyai Soc. Math. Stud.*, pages 199–262. Springer, Berlin, 2006.

[46] J. Lenz and D. Mubayi. Multicolor Ramsey Numbers for Complete Bipartite Versus Complete Graphs. submitted.

[47] J. Lenz and D. Mubayi. The poset of hypergraph quasirandomness. submitted.

[48] L. Lovász. Subgraph densities in signed graphons and the local Simonovits-Sidorenko conjecture. *Electron. J. Combin.*, 18:Paper 127, 21, 2011.

[49] L. Lovász and V. T. Sós. Generalized quasirandom graphs. *J. Combin. Theory Ser. B*, 98:146–163, 2008.

[50] L. Lu and X. Peng. High-ordered random walks and generalized laplacians on hypergraphs.

[51] L. Lu and X. Peng. Loose laplacian spectra of random hypergraphs.

[52] M. G. Martínez. The finite upper half space and related hypergraphs. *J. Number Theory*, 84:342–360, 2000.

[53] M. G. Martínez, H. M. Stark, and A. A. Terras. Some Ramanujan hypergraphs associated to $\mathrm{GL}(n, \mathbb{F}_q)$. *Proc. Amer. Math. Soc.*, 129:1623–1629 (electronic), 2001.

[54] J. S. Myers. Graphs without large complete minors are quasi-random. *Combin. Probab. Comput.*, 11:571–585, 2002.

[55] J. S. Myers and A. Thomason. The extremal function for noncomplete minors. *Combinatorica*, 25:725–753, 2005.

[56] B. Nagle, V. Rödl, and M. Schacht. The counting lemma for regular k-uniform hypergraphs. *Random Structures Algorithms*, 28:113–179, 2006.

[57] V. Nikiforov. Eigenvalues and extremal degrees of graphs. *Linear Algebra Appl.*, 419:735–738, 2006.

[58] V. Rödl and J. Skokan. Regularity lemma for k-uniform hypergraphs. *Random Structures Algorithms*, 25:1–42, 2004.

[59] V. Rödl and J. Skokan. Applications of the regularity lemma for uniform hypergraphs. *Random Structures Algorithms*, 28:180–194, 2006.

[60] A. Shapira. Quasi-randomness and the distribution of copies of a fixed graph. *Combinatorica*, 28:735–745, 2008.

[61] A. Shapira and R. Yuster. The effect of induced subgraphs on quasi-randomness. *Random Structures Algorithms*, 36:90–109, 2010.

[62] A. Shapira1 and R. Yuster. The quasi-randomness of hypergraph cut properties. *Random Structures Algorithms*, 40:105–131, 2012.

[63] A. Sidorenko. A correlation inequality for bipartite graphs. *Graphs Combin.*, 9:201–204, 1993.

[64] A. Sidorenko. An analytic approach to extremal problems for graphs and hypergraphs. In *Extremal problems for finite sets (Visegrád, 1991)*, volume 3 of *Bolyai Soc. Math. Stud.*, pages 423–455. János Bolyai Math. Soc., Budapest, 1994.

[65] M. Simonovits and V. T. Sós. Szemerédi's partition and quasirandomness. *Random Structures Algorithms*, 2:1–10, 1991.

[66] M. Simonovits and V. T. Sós. Hereditarily extended properties, quasi-random graphs and not necessarily induced subgraphs. *Combinatorica*, 17:577–596, 1997.

[67] M. Simonovits and V. T. Sós. Hereditary extended properties, quasi-random graphs and induced subgraphs. *Combin. Probab. Comput.*, 12:319–344, 2003.

[68] J. Skokan and L. Thoma. Bipartite subgraphs and quasi-randomness. *Graphs Combin.*, 20:255–262, 2004.

[69] C. K. Storm. The zeta function of a hypergraph. *Electron. J. Combin.*, 13:Research Paper 84, 26 pp. (electronic), 2006.

[70] T. Szabó. On the spectrum of projective norm-graphs. *Inform. Process. Lett.*, 86:71–74, 2003.

[71] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arith.*, 27:199–245, 1975.

[72] R. M. Tanner. Explicit concentrators from generalized N-gons. *SIAM J. Algebraic Discrete Methods*, 5:287–293, 1984.

[73] T. Tao. A variant of the hypergraph removal lemma. *J. Combin. Theory Ser. A*, 113:1257–1280, 2006.

[74] A. Thomason. Pseudorandom graphs. In *Random graphs '85 (Poznań, 1985)*, volume 144 of *North-Holland Math. Stud.*, pages 307–331. North-Holland, Amsterdam, 1987.

[75] A. Thomason. Random graphs, strongly regular graphs and pseudorandom graphs. In *Surveys in combinatorics 1987 (New Cross, 1987)*, volume 123 of *London Math. Soc. Lecture Note Ser.*, pages 173–195. Cambridge Univ. Press, Cambridge, 1987.

[76] R. Yuster. Quasi-randomness is determined by the distribution of copies of a fixed graph in equicardinal large sets. *Combinatorica*, 30:239–246, 2010.

A Algebraic properties of multilinear maps

In this section we prove several algebraic facts about multilinear maps, including Proposition 14 and Lemmas 15 and 16. Throughout this section, V and V_i are finite dimensional vector spaces over $\mathbb{R}$. Also in this section we make no distinction between bilinear maps and matrices, using whichever formulation is convenient. We will use a symbol $\cdot$ to denote the input to a linear map; for example, if $\phi : V_1 \times V_2 \times V_3 \to \mathbb{R}$ is a trilinear map and $x_1 \in V_1$ and $x_2 \in V_2$, then by the expression $\phi(x_1, x_2, \cdot)$ we mean the linear map from V_3 to $\mathbb{R}$ which takes a vector $x_3 \in V_3$ to $\phi(x_1, x_2, x_3)$. Lastly, we use several basic facts about tensors, all of which follow from the fact that we defined the tensor product of V and W to be the vector space over $\mathbb{R}$ of dimension $\dim(V)\dim(W)$. For example, if x and y are unit length, then $x \otimes y$ is also unit length.

Lemma 23. *Let $\phi : V \to \mathbb{R}$ be a linear map. There exists a vector v such that $\phi = \langle v, \cdot \rangle$.*

Proof. v is the vector dual to ϕ in the dual of the vector space V. Alternatively, let the ith coordinate of v be $\phi(e_i)$, since then for any x,

$$\phi(x) = \phi\left(\sum \langle x, e_i \rangle e_i\right) = \sum \langle x, e_i \rangle \phi(e_i) = \sum \langle x, e_i \rangle \langle v, e_i \rangle = \langle x, v \rangle .$$

□

Lemma 24. *Let $\phi, \psi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be t-linear maps. The maps $\phi * \psi$ and $A[\phi^{2^{t-1}}]$ are well defined. Also, $\phi * \psi$ is basis independent in the sense that the definition of $\phi * \psi$ is independent of the choice of orthonormal basis $b_1, \ldots, b_t$ of V_t.*

Proof. First, extending the definitions of $\phi * \psi$ and $A[\phi^{2^{t-1}}]$ linearly to the entire domain (non-simple tensors) is well defined, since ϕ and ψ are linear. That is, write each u_i and v_i in terms of some orthonormal basis and expand each tensor in $V_i \otimes V_i$ also in terms of this basis. The linearity of ϕ and ψ then shows that the definitions of $\phi * \psi$ and $A[\phi^{2^{t-1}}]$ are well defined and linear. To see basis independence of $\phi * \psi$, by Lemma 23 the linear map $\phi(u_1, \ldots, u_{t-1}, \cdot) : V_t \to \mathbb{R}$ equals $\langle u', \cdot \rangle$ for some vector u'. Similarly, $\psi(v_1, \ldots, v_t, \cdot)$ equals $\langle v', \cdot \rangle$ for some vector v'. Then

$$(\phi * \psi)(u_1 \otimes v_1, \ldots, u_{t-1} \otimes v_{t-1}) = \sum_{i=1}^{\dim(V_t)} \langle u', b_i \rangle \langle v', b_i \rangle = \langle u', v' \rangle .$$

The last equality is valid for any orthonormal basis, since the dot product of u' and v' sums the product of the ith coordinate of u' in the basis $\{b_1, \ldots, b_{\dim(V_t)}\}$ with the ith coordinate of v' in the basis $\{b_1, \ldots, b_{\dim(V_t)}\}$. □

Definition. *For $s \geq 0$ and V a finite dimensional vector space over $\mathbb{R}$, define the vector space isomorphism $\Gamma_{V,s} : V^{\otimes 2^s} \to V^{\otimes 2^s}$ as follows. If $s = 0$, define $\Gamma_{V,0}$ to be the identity map. If $s \geq 1$, let $\{b_1, \ldots, b_{\dim(V)}\}$ be any orthonormal basis of V and define for all $(i_1, \ldots, i_{2^{s-1}}, j_1, \ldots, j_{2^{s-1}}) \in [\dim(V)]^{2^s}$,*

$$\Gamma_{V,s}(b_{i_1} \otimes b_{j_1} \otimes \cdots \otimes b_{i_{2^{s-1}}} \otimes b_{j_{2^{s-1}}}) = b_{j_1} \otimes b_{i_1} \otimes \cdots \otimes b_{j_{2^{s-1}}} \otimes b_{i_{2^{s-1}}}. \tag{25}$$

Extend $\Gamma_{V,s}$ linearly to all of $V^{\otimes 2^s}$.

Remarks. $\Gamma_{V,s}$ is a vector space isomorphism since it restricts to a bijection of an orthonormal basis to itself. Also, it is easy to see that $\Gamma_{V,s}$ is well defined and independent of the choice of orthonormal basis, since each b_i can be written as a linear combination of an orthonormal basis $\{b'_1, \ldots, b'_{\dim(V)}\}$ and (25) can be expanded using linearity. For notational convenience, we will usually drop the subscript V and write Γ_s for $\Gamma_{V,s}$.

Lemma 25. *Let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map, let $0 \leq s \leq t-1$, and let $x_1 \in V_1^{\otimes 2^s}, \ldots, x_{t-s} \in V_{t-s}^{\otimes 2^s}$. Then*

$$\phi^{2^s}(x_1, \ldots, x_{t-s}) = \phi^{2^s}(\Gamma_s(x_1), \ldots, \Gamma_s(x_{t-s})).$$

Proof. By induction on s. The base case is $s = 0$ where Γ_0 is the identity map. Expand the definition of $\phi^{2^{s+1}}$ and use induction to obtain

$$\begin{aligned}\phi^{2^{s+1}}(x_1 \otimes y_1, \ldots, x_{t-s-1} \otimes y_{t-s-1}) &= \sum_{j=1}^{\dim(V_{t-s}^{\otimes 2^s})} \phi^{2^s}(x_1, \ldots, x_{t-s-1}, b_j)\phi^{2^s}(y_1, \ldots, y_{t-s-1}, b_j) \\ &= \sum_{j=1}^{\dim(V_{t-s}^{\otimes 2^s})} \phi^{2^s}\big(\Gamma_s(x_1), \ldots, \Gamma_s(x_{t-s-1}), \Gamma_s(b_j)\big)\phi^{2^s}\big(\Gamma_s(y_1), \ldots, \Gamma_s(y_{t-s-1}), \Gamma_s(b_j)\big).\end{aligned}$$

But since Γ_s is a vector space isomorphism, $\{\Gamma_s(b_1), \ldots, \Gamma_s(b_{\dim(V_{t-s}^{\otimes 2^s})})\}$ is an orthonormal basis of $V_{t-s}^{\otimes 2^s}$. Thus Lemma 24 shows that

$$\sum_{j=1}^{\dim(V_{t-s}^{\otimes 2^s})} \phi^{2^s}\big(\Gamma_s(x_1), \ldots, \Gamma_s(x_{t-s-1}), \Gamma_s(b_j)\big)\phi^{2^s}\big(\Gamma_s(y_1), \ldots, \Gamma_s(y_{t-s-1}), \Gamma_s(b_j)\big)$$
$$= \phi^{2^{s+1}}\big(\Gamma_s(x_1) \otimes \Gamma_s(y_1), \ldots, \Gamma_s(x_{t-s-1}) \otimes \Gamma_s(y_{t-s-1})\big)$$

Finally, $\Gamma_s(x_i) \otimes \Gamma_s(y_i) = \Gamma_{s+1}(x_i \otimes y_i)$ (write x_i and y_i as linear combinations, expand $\Gamma_{s+1}(x_i \otimes y_i)$ using linearity, and apply (25)). Thus $\phi^{2^{s+1}}(x_1 \otimes y_1, \ldots, x_{t-s-1} \otimes y_{t-s-1}) = \phi^{2^{s+1}}(\Gamma_{s+1}(x_1 \otimes y_1), \ldots, \Gamma_{s+1}(x_{t-s-1} \otimes y_{t-s-1}))$, completing the proof. □

Proof of Lemma 16. Let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map. $A[\phi^{2^{t-1}}]$ is a bilinear map from $V_1^{\otimes 2^{t-2}} \times V_1^{\otimes 2^{t-2}} \to \mathbb{R}$ and so is a square matrix of dimension $\dim(V_1)^{2^{t-2}}$. Lemma 25 shows that $A[\phi^{2^{t-1}}]$ is a symmetric matrix, since

$$\begin{aligned} A[\phi^{2^{t-1}}](x_1 \otimes \cdots \otimes x_{2^{t-2}}, y_1 \otimes \cdots \otimes y_{2^{t-2}}) &= \phi^{2^{t-1}}(x_1 \otimes y_1 \otimes \cdots \otimes x_{2^{t-2}} \otimes y_{2^{t-2}}) \\ &= \phi^{2^{t-1}}(\Gamma(x_1 \otimes y_1 \otimes \cdots \otimes x_{2^{t-2}} \otimes y_{2^{t-2}})) \\ &= \phi^{2^{t-1}}(y_1 \otimes x_1 \otimes \cdots \otimes y_{2^{t-2}} \otimes x_{2^{t-2}}) \\ &= A[\phi^{2^{t-1}}](y_1 \otimes \cdots \otimes y_{2^{t-2}}, x_1 \otimes \cdots \otimes x_{2^{t-2}}). \end{aligned}$$

The above equation is valid for all $x_i, y_i \in V_1$, in particular for all basis elements of V_1 which implies that $A[\phi^{2^{t-1}}](w, z) = A[\phi^{2^{t-1}}](z, w)$ for all basis vectors w, z of $V_1^{\otimes 2^{t-2}}$. Thus $A[\phi^{2^{t-1}}]$ is a square symmetric real-valued matrix. □

Lemma 26. *Let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map and let $x_1 \in V_1, \ldots, x_t \in V_t$ be unit length vectors. Then*

$$|\phi(x_1, \ldots, x_t)|^2 \leq \left|\phi^2(x_1 \otimes x_1, \ldots, x_{t-1} \otimes x_{t-1})\right|.$$

Proof. Consider the linear map $\phi(x_1, \ldots, x_{t-1}, \cdot)$ which is a linear map from V_t to $\mathbb{R}$. By Lemma 23, there exists a vector $w \in V_t$ such that $\phi(x_1, \ldots, x_{t-1}, \cdot) = \langle w, \cdot \rangle$. Now expand out the definition of ϕ^2:

$$\phi^2(x_1 \otimes x_1, \ldots, x_{t-1} \otimes x_{t-1}) = \sum_j |\phi(x_1, \ldots, x_{t-1}, b_j)|^2 = \sum_j |\langle w, b_j \rangle|^2 = \langle w, w \rangle$$

where the last equality is because $\{b_j\}$ is an orthonormal basis of V_t. Since $\|w\| = \sqrt{\langle w, w \rangle}$,

$$\left|\phi^2(x_1 \otimes x_1, \ldots, x_{t-1} \otimes x_{t-1})\right| = |\langle w, w \rangle| = \left|\left\langle w, \frac{w}{\|w\|} \right\rangle\right|^2.$$

But since x_t is unit length and $\langle w, \cdot \rangle$ is maximized over the unit ball at vectors parallel to w (so maximized at $w/\|w\|$), $\left|\left\langle w, \frac{w}{\|w\|}\right\rangle\right| \geq |\langle w, x_t\rangle|$. Thus

$$\left|\phi^2(x_1 \otimes x_1, \ldots, x_{t-1} \otimes x_{t-1})\right| = \left|\left\langle w, \frac{w}{\|w\|}\right\rangle\right|^2 \geq |\langle w, x_t\rangle|^2 = |\phi(x_1, \ldots, x_t)|^2 .$$

The last equality used the definition of w, that $\phi(x_1, \ldots, x_{t-1}, \cdot) = \langle w, \cdot \rangle$. □

Lemma 27. *Let $\phi : V_1 \times \cdots \times V_t \to \mathbb{R}$ be a t-linear map and let $x_1 \in V_1, \ldots, x_t \in V_t$ be unit length vectors. Then for $0 \leq s \leq t-1$,*

$$|\phi(x_1, \ldots, x_t)|^{2^s} \leq \left|\phi^{2^s}(\underbrace{x_1 \otimes \cdots \otimes x_1}_{2^s}, \ldots, \underbrace{x_{t-s} \otimes \cdots \otimes x_{t-s}}_{2^s})\right|$$

which implies that

$$|\phi(x_1, \ldots, x_t)|^{2^{t-1}} \leq \left|A[\phi^{2^{t-1}}](\underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-2}}, \underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-2}})\right| .$$

Proof. By induction on s. The base case is $s = 0$ where both sides are equal and the induction step follows from Lemma 26. By definition of $A[\phi^{2^{t-1}}]$,

$$\left|A[\phi^{2^{t-1}}](\underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-2}}, \underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-2}})\right| = \left|\phi^{2^{t-1}}(\underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-1}})\right| ,$$

completing the proof. □

Proof of Lemma 15. Pick $x_1, \ldots, x_t$ unit length vectors to maximize ϕ, so $\phi(x_1, \ldots, x_t) = \|\phi\|$. Then Lemma 27 shows that

$$\|\phi\|^{2^{t-1}} = |\phi(x_1, \ldots, x_t)|^{2^{t-1}} \leq \left|A[\phi^{2^{t-1}}](\underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-2}}, \underbrace{x_1 \otimes \cdots \otimes x_1}_{2^{t-2}})\right|$$

Since $x_1 \otimes \cdots \otimes x_1$ is unit length, the above expression is upper bounded by the spectral norm of $A[\phi^{2^{t-1}}]$. □

Lemma 28. *Let $\{M_r\}_{r\to\infty}$ be a sequence of square symmetric real-valued matrices with dimension going to infinity where $\lambda_2(M_r) = o(\lambda_1(M_r))$. Let u_r be a unit length eigenvector corresponding to the largest eigenvalue in absolute value of M_r. If $\{x_r\}$ is a sequence of unit length vectors such that $\left|x_r^T M_r x_r\right| = (1 + o(1))\lambda_1(M_r)$, then*

$$\|u_r - x_r\| = o(1).$$

Consequently, for any unit length sequence $\{y_r\}$ where each y_r is perpendicular to x_r,

$$\left|y_r^T M_r y_r\right| = o(\lambda_1(M_r)).$$

Proof. Throughout this proof, the subscript r is dropped; all terms $o(\cdot)$ should be interpreted as $r \to \infty$. This exact statement was proved by Chung, Graham, and Wilson [18], although they don't clearly state it as such. We give a proof here for completeness using slightly different language but the same proof idea: if x projected onto $u^\perp$ is too big then the second largest eigenvalue is too big. Write $x = \alpha v + \beta u$ where v is a unit length vector perpendicular to u and $\alpha, \beta \in \mathbb{C}$ and $\alpha^2 + \beta^2 = 1$ (since u is an eigenvector it might have complex entries). Let $\phi(x, y) = x^T M y$ be the bilinear map corresponding to M. Since $u^T M v = \lambda_1 u^T v = \lambda_1 \langle u, v \rangle = 0$, we have $\phi(u, v) = 0$. This implies that

$$\begin{aligned}\phi(x, x) &= \phi(\alpha v + \beta u, \alpha v + \beta u) = \alpha^2 \phi(v, v) + \beta^2 \phi(u, u) + 2\alpha\beta\phi(u, v) \\ &= \alpha^2 \phi(v, v) + \beta^2 \phi(u, u).\end{aligned}$$

The second largest eigenvalue of M is the largest eigenvalue of $M - \lambda_1(M) u u^T$ which is the spectral norm of $M - \lambda_1(M) u u^T$. Thus

$$|\phi(v, v)| = |v^T M v| = |v^T (M - \lambda_1(M) u u^T) v| \leq \lambda_2(M). \tag{26}$$

Using that $\phi(u, u) = \lambda_1(M)$ and the triangle inequality, we obtain

$$|\phi(x, x)| \leq \alpha^2 \lambda_2(M) + \beta^2 \lambda_1(M). \tag{27}$$

Since $\alpha^2 + \beta^2 = 1$, $|\alpha|$ and $|\beta|$ are between zero and one. Combining this with (27) and $|\phi(x, x)| = (1 + o(1))\lambda_1(M)$ and $\lambda_2(M) = o(\lambda_1(M))$, we must have $|\beta| = 1 + o(1)$ which in turn implies that $|\alpha| = o(1)$. Consequently,

$$\|u - x\|^2 = \langle u - x, u - x \rangle = \langle u, u \rangle + \langle x, x \rangle - 2 \langle u, x \rangle = 2 - 2\beta = o(1).$$

Now consider some y perpendicular to x and similarly to the above, write $y = \gamma w + \delta u$ for some unit length vector w perpendicular to u and $\gamma, \delta \in \mathbb{C}$ with $\gamma^2 + \delta^2 = 1$. Then

$$\phi(y, y) = \phi(\gamma w + \delta u, \gamma w + \delta u) = \gamma^2 \phi(w, w) + \delta^2 \phi(u, u)$$

and as in (26), we have $|\phi(w, w)| \leq \lambda_2(M)$. Thus

$$|\phi(y, y)| \leq \gamma^2 \lambda_2(M) + \delta^2 \lambda_1(M).$$

We want to conclude that the above expression is $o(\lambda_1(M))$. Since $\lambda_2(M) = o(\lambda_1(M))$, we must prove that $|\delta| = o(1)$ to complete the proof.

$$\delta = \langle y, u \rangle = \left\langle y, \frac{x - \alpha v}{\beta} \right\rangle = \frac{1}{\beta} \big(\langle y, x \rangle - \alpha \langle y, v \rangle \big) = \frac{-\alpha \langle y, v \rangle}{\beta}.$$

But $|\alpha| = o(1)$, $|\beta| = 1 + o(1)$, and $\|y\| = \|v\| = 1$ so $|\delta| = o(1)$ as required. □

Lemma 29. *Let $J : V_1 \times \cdots \times V_t \to \mathbb{R}$ be the all-ones map and let $\vec{1}_i$ be the all-ones vector in V_i. Then for all $x_1, \ldots, x_t$ with $x_i \in V_i$,*

$$J(x_1, \ldots, x_t) = \left\langle \vec{1}_1, x_1 \right\rangle \cdots \left\langle \vec{1}_t, x_t \right\rangle. \tag{28}$$

Proof. If $x_1, \ldots, x_t$ are standard basis vectors, then the left and right hand side of (28) are the same. By linearity, (28) is then the same for all $x_1, \ldots, x_t$. □

Proof of Proposition 14. Again throughout this proof, the subscript r is dropped; all terms $o(\cdot)$ should be interpreted as $r \to \infty$. Let $\hat{1}$ denote the all-ones vector scaled to unit length in the appropriate vector space. Pick an ordering $\vec{\pi} = (k_1, \ldots, k_t)$ of π. The definition of spectral norm is independent of the choice of the ordering for the entries of $\vec{\pi}$, so $\|\psi_{\vec{\pi}} - qJ_{\vec{\pi}}\|$ is the same for all orderings. Let $w_1, \ldots, w_t$ be unit length vectors where $(\psi_{\vec{\pi}} - qJ_{\vec{\pi}})(w_1, \ldots, w_t) = \|\psi_{\vec{\pi}} - qJ_{\vec{\pi}}\|$ and write $w_i = \alpha_i y_i + \beta_i \hat{1}$ where y_i is a unit length vector perpendicular to the all-ones vector and $\alpha_i, \beta_i \in \mathbb{R}$ with $\alpha_i^2 + \beta_i^2 = 1$. Then

$$\begin{aligned}\|\psi_{\vec{\pi}} - qJ_{\vec{\pi}}\| &= (\psi_{\vec{\pi}} - qJ_{\vec{\pi}})(w_1, \ldots, w_t) = (\psi_{\vec{\pi}} - qJ_{\vec{\pi}})(\alpha_1 y_1 + \beta_1 \hat{1}, \ldots, \alpha_t y_t + \beta_t \hat{1}) \\ &= \psi_{\vec{\pi}}(\alpha_1 y_1 + \beta_1 \hat{1}, \ldots, \alpha_t y_t + \beta_t \hat{1}) - q \dim(V_r)^{k/2} \prod \beta_i. \end{aligned} \tag{29}$$

The last equality used that y_i is perpendicular to $\hat{1}$, so Lemma 29 implies that if y_i appears as input to $J_{\vec{\pi}}$ then the outcome is zero no matter what the other vectors are. Thus the only non-zero term involving $J_{\vec{\pi}}$ is $J_{\vec{\pi}}(\hat{1}, \ldots, \hat{1}) = \dim(V_r)^{k/2}$. Note that $\psi(\hat{1}, \ldots, \hat{1}) = \psi_{\vec{\pi}}(\hat{1}, \ldots, \hat{1})$ since the all-ones vector scaled to unit length in $V^{\otimes k_i}$ is the tensor product of the all-ones vector scaled to unit length in V. Inserting $q = \dim(V_r)^{-k/2} \psi_{\vec{\pi}}(\hat{1}, \ldots, \hat{1})$ in (29), we obtain

$$\|\psi_{\vec{\pi}} - qJ_{\vec{\pi}}\| = \psi_{\vec{\pi}}(\alpha_1 y_1 + \beta_1 \hat{1}, \ldots, \alpha_t y_t + \beta_t \hat{1}) - \left(\prod \beta_i\right) \psi_{\vec{\pi}}(\hat{1}, \ldots, \hat{1}). \tag{30}$$

Now consider expanding $\psi_{\vec{\pi}}$ in (30) using linearity; the term $(\prod \beta_i)\psi_{\vec{\pi}}(\hat{1}, \ldots, \hat{1})$ cancels, so all terms include at least one y_i. We claim that each of these terms is small; the following claim finishes the proof, since $\|\psi_{\vec{\pi}} - qJ_{\vec{\pi}}\|$ is the sum of terms each of which $o(\psi(\hat{1}, \ldots, \hat{1}))$.

Claim: If $z_1, \ldots, z_{i-1}, z_{i+1}, \ldots, z_t$ are unit length vectors, then

$$|\psi_{\vec{\pi}}(z_1, \ldots, z_{i-1}, y_i, z_{i+1}, \ldots, z_t)| = o(\psi(\hat{1}, \ldots, \hat{1})).$$

Proof. Change the ordering of $\vec{\pi}$ to an ordering $\vec{\pi}'$ that differs from $\vec{\pi}$ by swapping 1 and i. Since ψ is symmetric,

$$\psi_{\vec{\pi}}(z_1, \ldots, z_{i-1}, y_i, z_{i+1}, \ldots, z_t) = \psi_{\vec{\pi}'}(y_i, z_2, \ldots, z_{i-1}, z_1, z_{i+1}, \ldots, z_t). \tag{31}$$

Therefore proving the claim comes down to bounding $\psi_{\vec{\pi}'}(y_i, z_2, \ldots, z_{i-1}, z_1, z_{i+1}, \ldots, z_t)$, which is a combination of Lemma 27 and Lemma 28 as follows. For the remainder of this proof, denote by A the matrix $A[\psi_{\vec{\pi}'}^{2^{t-1}}]$. By assumption, we have $\lambda_2(A) = o(\lambda_1(A))$ so Lemma 28 can be applied to the matrix sequence A. Next we would like to show that we can use $\hat{1}$ for x in the statement of Lemma 28; i.e. that $A(\hat{1}, \hat{1}) = (1 + o(1))\lambda_1(A)$. By Lemma 27 and the assumption $\lambda_1(A) = (1 + o(1))\psi(\hat{1}, \ldots, \hat{1})^{2^{t-1}}$, we have

$$\left|\psi_{\vec{\pi}'}(\hat{1}, \ldots, \hat{1})\right|^{2^{t-1}} \leq \left|A(\hat{1}, \hat{1})\right| \leq \lambda_1(A) = (1 + o(1))\psi\left(\hat{1}, \ldots, \hat{1}\right)^{2^{t-1}}.$$

Using the definition of $\psi_{\vec{\pi}'}$, we have $\psi_{\vec{\pi}'}(\hat{1},\ldots,\hat{1}) = \psi(\hat{1},\ldots,\hat{1})$, which implies asymtotic equality through the above equation. In particular, $|A(\hat{1},\hat{1})| = (1+o(1))\lambda_1(A)$ which is the condition in Lemma 28 for $x = \hat{1}$. Lastly, to apply Lemma 28 we need a vector y perpendicular to $\hat{1}$. The vector $y_i \otimes \cdots \otimes y_i \in V^{\otimes k_i 2^{t-2}}$ is pependicular to $\hat{1}$ (in $V^{\otimes k_i 2^{t-2}}$) since y_i itself is perpendicular to $\hat{1}$ (in $V^{\otimes k_i}$). Thus Lemma 28 implies that

$$\left| A(\underbrace{y_i \otimes \cdots \otimes y_i}_{2^{t-2}}, \underbrace{y_i \otimes \cdots \otimes y_i}_{2^{t-2}}) \right| = o(\lambda_1(A)). \tag{32}$$

Using Lemma 27 again shows that

$$|\psi_{\vec{\pi}'}(y_i, z_2, \ldots, z_{i-1}, z_1, z_{i+1}, \ldots, z_t)|^{2^{t-1}} \leq \left| A(\underbrace{y_i \otimes \cdots \otimes y_i}_{2^{t-2}}, \underbrace{y_i \otimes \cdots \otimes y_i}_{2^{t-2}}) \right|.$$

Combining this equation with (31) and (32) shows that $|\psi_{\vec{\pi}}(z_1, \ldots, z_{i-1}, y_i, z_{i+1}, \ldots, z_t)|^{2^{t-1}} = o(\lambda_1(A))$. By assumption, $\lambda_1(A) = (1+o(1))\psi(\hat{1},\ldots,\hat{1})^{2^{t-1}}$, completing the proof of the claim.

□